

TIBCO BusinessEvents®

Developer's Guide

Software Release 5.1.2

February 2014

Document Updated: July 2014

Important Information

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENT IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE "LICENSE" FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document contains confidential information that is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of TIBCO Software Inc.

TIBCO, The Power of Now, TIBCO ActiveMatrix, TIBCO ActiveMatrix BusinessWorks, TIBCO Administrator, TIBCO ActiveSpaces, TIBCO Designer, TIBCO Enterprise Message Service, TIBCO Hawk, TIBCO Runtime Agent, TIBCO Rendezvous, are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

EJB, Java EE, J2EE, and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

THIS SOFTWARE MAY BE AVAILABLE ON MULTIPLE OPERATING SYSTEMS. HOWEVER, NOT ALL OPERATING SYSTEM PLATFORMS FOR A SPECIFIC SOFTWARE VERSION ARE RELEASED AT THE SAME TIME. SEE THE README.TXT FILE FOR THE AVAILABILITY OF THIS SOFTWARE VERSION ON A SPECIFIC OPERATING SYSTEM PLATFORM.

THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "READ ME" FILES.

This product is covered by U.S. Patent No. 7,472,101.

Copyright © 2004-2014 TIBCO Software Inc. ALL RIGHTS RESERVED.

TIBCO Software Inc. Confidential Information

Contents

Tables	xxi
Preface	xxv
Changes from the Previous Release of this Guide	xxvi
TIBCO BusinessEvents Express	xxxii
Related Documentation	xxxii
TIBCO BusinessEvents and Add-On Product Documentation	xxxii
Accessing TIBCO BusinessEvents Functions Reference Documentation	xxxvi
Other TIBCO Product Documentation	xxxvi
Typographical Conventions	xxxvii
Connecting with TIBCO Resources	xl
How to Join TIBCOCommunity	xl
How to Access TIBCO Documentation	xl
How to Contact TIBCO Support	xl
Chapter 1 Project Tasks	1
Project Development Overview	2
Creating a Project	3
Finding a Project Element	4
Exporting (Generating) Concept and Event Schema (XSD) Files	5
Validating a Project or Project Resource	7
Working with External Library and Custom Function Paths	8
Working with Project Libraries	10
Creating (Exporting) a Project Library	10
Adding (and Removing) Project Libraries in a Project	12
Working with Global Variables	14
Setting and Overriding Global Variables from a Project Library	14
Adding and Managing Global Variables	15
Using Global Variables	18
Overriding Global Variables at Deploy Time	18
Order of Precedence of Global Variable Overrides	19
Storing Trusted Certificates Outside of Your Project	20
Building an Enterprise Archive (EAR File)	22
EAR File Encoding	22

Options for Deploying the EAR	22
Building an EAR File in TIBCO BusinessEvents Studio	23
Chapter 2 TIBCO BusinessEvents Studio Tools Utility	25
Building an Enterprise Archive (EAR File) at the Command Line	26
Importing an Older Studio Project at the Command Line	28
Working With Project Libraries at the Command Line	31
Migrating Core Coherence Functions at the Command Line	34
Generating All Project Class Files at the Command Line	35
Generating Encrypted Passwords	37
Running the Tester on the Command Line	38
Command Line Options	39
generate	40
import	41
assert	42
Chapter 3 Element Refactoring Operations	43
Renaming, Moving, Deleting, and Copy-Pasting Elements	44
Project Level Actions	45
Renaming, Moving, and Deleting Elements	45
Copy-Pasting an Element	49
Migrating Core Coherence Functions	50
Automatic Refactoring Actions and Limitations	53
Refactoring Limitations	53
Refactoring for Move and Rename Operations	53
Refactoring for Delete Operations	55
Chapter 4 Channels and Destinations	57
Overview of Channels and Destinations	58
Types of Channels	58
Selecting a Serializer	60
Deploy-time Configuration	60
Mapping Incoming Messages to Non-default Events	61
Working with Rendezvous Channels	62
Rendezvous Message Header	62
Basic Serializer	62
Serializer For Use with Payloads	62
Avoiding REGISTRATION COLLISION RVCM Advisory Messages	63
Working with Local Channels	65
Event Use Count	65

Using Local Channel Destinations	65
Adding Channels and Destinations	66
Adding a Destination to a Channel	68
Communicating with Other Sources using TCP	70
Example TCP Rule Function to Start a Local TCP Server	70
Example TCP Rule Function to Connect to a Remote TCP Server	71
APIs for TCP Communication	71
Channel Resource Reference	73
Wizard and Configuration Section	73
Destination Resource Reference	78
Chapter 5 JMS Channels	83
Overview of JMS Channels	84
Selecting a JMS Serializer	85
BytesMessageSerializer	86
UtfBytesMessageSerializer	86
TextMessageSerializer	86
MessageWithNoBody	86
Creating Unique JMS DurableSubscriber Name Properties	87
Setting the JMS Message Acknowledgement Mode	89
Using CLIENT_ACKNOWLEDGE Mode with Websphere MQ and Cache-Aside	91
When JMS Messages are Acknowledged	92
Using JMS Header Properties in Incoming and Outgoing Messages	93
Setting Certain Header Properties in Destinations	93
Setting Header Properties Using Header Properties from Incoming JMS Messages	93
Setting JMS Header Properties in Outgoing JMS Messages Using Event Properties	93
How TIBCO BusinessEvents Sets the JMSReplyTo Header in Outgoing Messages	94
JMS Header Field Names	95
Chapter 6 HTTP and SOAP Channels	97
Overview of HTTP and SOAP Channels	98
SOAP Support	98
Working with HTTP Requests	99
Event Based Approach	99
Action Rule Function Based Approach	100
Configuring TIBCO BusinessEvents to Receive and Send HTTP Requests	102
Setting Up Fault Tolerance for the HTTP Channel	106
HTTP Channel Advanced Configuration Settings	108
Defining Event Properties for Standard HTTP Header Properties	114
Using HTTP Functions to Configure HTTP Request Messages	116

Generating a Self-Signed SSL Certificate (Keystore)	116
Getting POST Data	117
Loading Trusted Certificates	117
Sending an Event	118
Sending an Asynchronous Request (Not Secure)	118
Sending a Secure Asynchronous Request	120
Sending a Secure Synchronous Request	120
Configuring TIBCO BusinessEvents as a SOAP Server and Client	122
Overview of SOAP Related Resources	122
Mapping of SOAP Request URI to Destination	123
Manually Creating Resources to Work with SOAP Services	124
Creating Resources Using the WSDL Import Utility	125
Exporting Resources as a WSDL File	127
Parsing and Building SOAP Messages	130
Working with Incoming SOAP Messages (Event Payloads)	130
Working with Outgoing SOAP Messages (Event Payloads)	133
Understanding the WSDL to Project Resource Mapping	135
Example WSDL	135
Example Project Folder Structure	136
How Project Artifacts are Named	138
Chapter 7 HAWK Channel	141
Working with the Hawk Channel	142
Runtime Configuration for the Hawk Channel	143
Using the Hawk Destination and Event Wizard	144
Creating a Destination and Event Using the Wizard	144
Create Event for Hawk Catalog Function Using Wizard	145
Hawk Channel Destination Reference	146
Chapter 8 ActiveSpaces Channel	148
Basic TIBCO ActiveSpaces Concepts and Terminology	149
Working with ActiveSpaces Channels	155
Creating a Destination and Event Using the Wizard	155
Configuring the Destination and Default Event Manually	158
Support for ActiveSpaces Security	159
Chapter 9 Simple Events	161
Overview of Simple Events	162
Using Inheritance	162
Working with Events in Rules	164
Explicitly Assert Events Created in Rules	164

Specifying Default and Non-Default Destinations	164
Scheduling Simple Events	164
Adding a Simple Event	165
Defining Payloads Using XML Schema Files	165
Simple Event Reference	167
Wizard and Configuration (Standard Tab)	167
Properties (Standard Tab)	169
Declaration and Expiry Action (Advanced Tab)	169
Payload (Advanced Tab)	170
Simple Event Attributes Reference	173
Chapter 10 Time Events and Scheduler Functions	175
Overview of Time Events	176
Scheduled Time Events	176
Rule Based Time Events	176
Working With Time Events	178
Adding a Time Event	178
Configuring a Rule Based Time Event in a Rule or Rule Function	178
TimeEvent Resource Reference	180
Wizard and Configuration Tab	180
TimeEvent Attributes Reference	182
Rule Based TimeEvent Function Reference	184
Using Scheduler Functions (Requires Cache OM)	185
Chapter 11 Advisory Events	187
Working With Advisory Events	188
Uses of Advisory Events	188
Adding an Advisory Event to a Rule	188
Advisory Event Attributes Reference	190
Chapter 12 Concepts	193
Overview of Concepts	194
Adding Concepts and Concept Relationships	195
Adding a Concept	195
Adding Concept Relationships	195
Concept Resource Reference	197
Wizard and Configuration Tab	197
Properties	198
Metadata	199
Concept Attributes Reference	200

Chapter 13 Scorecards	203
Understanding and Working With Scorecards	204
Adding a Scorecard	204
Using a Scorecard in Rules	205
Scorecard Resource Reference	205
Chapter 14 Domain Models	207
Overview of Domain Models	208
Domain Model Value Descriptions for User Friendly Presentation	208
Adding a Domain Model	209
Importing and Exporting Domain Models	212
Export To and Import From Excel	212
Import From a Database Table	213
Associating Domain Models with Properties	215
Validating Data in Domain Models	216
Chapter 15 Shared Resources	217
Adding a Shared Resource	218
Enabling the Test Connection Feature	219
ActiveSpaces Connection	220
Wizard and Configuration Tab	220
Hawk Connection	223
Wizard and Configuration Tab	223
HTTP Connection	225
Wizard and Configuration Tab	225
Configure SSL	226
Identity Resource	228
Wizard and Configuration Tab	228
JDBC Connection	230
Wizard and Configuration Tab	230
Connection Pooling	233
Test Connection Button	233
JMS Application Properties	234
Wizard and Configuration Tab	234
JMS Connection	235
Wizard and Configuration Tab	235
Test Connection Button	238
Advanced Tab	238
Configure SSL	239
JNDI Configuration	242

Wizard and Configuration Tab	242
Advanced Section	243
Rendezvous Transport	245
Wizard and Configuration Tab	245
Configure SSL Button	246
Advanced Section	247
Chapter 16 Rules and Rule Functions	251
Overview of Rules and Rule Functions	252
Form-based and Source Rule Editors	252
Rule Components	252
Effect of Cache Only Cache Mode	253
Adding a Rule	254
Rule Editor Reference	258
Adding a Rule Function	260
Rule Function Resource Reference	263
Using Variables and Functions in the Rule Editor	265
Using Catalog Functions in the Rule Editor	265
Using Global Variables in the Rule Editor	265
Using the Function Argument Mapper	266
Using Priority and Rank to Control Order of Rule Execution	269
Using the Quick Fix Feature in the Rule Editor	271
Tips for Working in the Rule Editor	273
Event Preprocessors	275
Configuring an Event Preprocessor	275
Transaction Error Handler Rule Function	276
Chapter 17 Rule Templates	279
Rule Template and Rule Template View Overview	280
The Rule Template Editor	280
Rule Template Views	283
Adding a Rule Template	284
Adding a Rule Template View	286
Rule Template Editor Reference	287
Chapter 18 Functions	289
Overview of Catalog Functions	290
Built-in Functions	291
Standard Functions	291

ActiveSpaces Functions	293
Process Orchestration Functions	293
BusinessWorks Functions	293
CEP Load Balancer Functions	293
CEP Pattern Functions	294
CEP Query Functions	294
Communication Functions	294
Hawk Functions	294
DBMS Functions	294
Security Functions	294
Studio Functions	294
Custom Functions	295
Adding Custom Functions	295
Task Summary	295
Restrictions on Use of Custom Functions	296
Example Custom Function with Tooltips	297
Ontology Functions	299
Enabling Extended Functions	300
Function Tooltips and Decorations	301
Tool Tips	301
Decorations Indicating Where Functions can be Called	301
Temporal Functions and Their Parameters	303
Virtual Rule Functions and VRF Catalog Functions	305
Cache Related Functions	309
Indexing for More Efficient Cache Queries	311
Enabling Explicit Tuple Format for DataGrid Fields	312
Example Rule Function for a TIBCO BusinessEvents DataGrid Cache	314
Oracle Coherence Cache Query Functions	315
Constants, Extractors, and Filters Categories	315
Query Category	315
Structure of a Function Catalog	317
Example Function Catalog	317
Elements in the Function Catalog	318
Using the Reevaluate Element	319
Chapter 19 Rule Language Grammar	321
Rule Language Basics	322
Whitespace	322
Comments	322
Separators	323
Identifier Naming Requirements	323
Local Variables	324

Literals	325
Escape Sequences	326
Operators	326
Keywords and Other Reserved Words	328
Attributes	329
Working with Concept and Event Properties	331
Accessing a Concept Property Atom	331
Working with Concept Property Arrays	332
Event Property	333
Exception Handling	334
Syntax	334
Examples	335
Flow Control	337
if/else	337
for	337
while	338
Chapter 20 Rule Language Datatypes	339
Concept Properties to XML Datatype Conversions	340
Compatibility of Operators with Types	341
Correcting Inconsistencies of Type	343
String Operands	343
Arithmetic Expressions	343
Assignment Conversion	344
Function Argument Conversion	344
Chapter 21 Mapping and Transforming Data	345
Overview of Mapping and Transformation	346
Function Section	346
Input Section	346
Mapping and Transforming Data to Function Input	347
Statements, Hints, and Errors	347
Buttons, Menus, and Icons	349
Toolbar and Right-Click Menu on the Input Section	349
Icons for Schema Element Datatypes	352
Qualifier Icons	353
Specifying Constants	355
Date and Datetime Strings in Constants	355
Data Validation	356
Repairing Incorrect Mappings	357
Shortcuts	358

Statement Menu Options	358
Dragging to the Left	358
Cutting and Pasting	359
Automatic Testing (at Runtime)	360
Examples of Mappings	362
Setting an Element Explicitly to Nil	362
Merging Input from Multiple Sources	363
Converting a List Into a Grouped List	367
Merging Two Corresponding Lists	371
Coercions	374
XSLT Statements	378
Attribute	378
Choose	378
Comment	379
Copy	379
Copy-Contents-Of	379
Copy-Of	379
Element	380
For-Each	380
For-Each-Group	380
Generate Comment	380
Generate PI	381
If	381
Value-Of	381
Variable	381
Chapter 22 XPath Formula Builder	383
XPath Basics	384
Addressing Schema Elements	384
Evaluation Context	385
Namespaces	385
Search Predicates	385
Testing for Nil	386
Comments	386
The XPath Formula Builder	387
String Representations of Datatypes	390
Date and Time Functions	391
Chapter 23 Cluster Deployment Descriptor (CDD)	395
Cluster Deployment Descriptor Overview	396
Maintaining the CDD File	396
Defining and Configuring the Cluster and Object Management Type	397

Configuring Management of Domain (Entity) Object Instances	397
Defining and Configuring Agent Classes and Processing Units	398
How CDD Settings Apply at Runtime	400
Using Properties at Different Levels	400
Order of Precedence at Runtime	400
Setting Global Variables in the CDD File (for Command Line Startup)	402
Adding a Cluster Deployment Descriptor (CDD)	403
CDD Cluster Tab General Settings Reference	405
Chapter 24 Cache OM and Cluster Configuration	407
Cache Object Management Configuration	408
Configuring a Limited (or Unlimited) Cache	408
Ensuring Multiple Clusters Do Not Conflict	410
CDD Cluster Tab and Cache OM Settings Reference	411
Synchronous and Asynchronous Replication of Cache Objects	414
Property for Cache Based Object Management on AIX	415
Configuring Cluster Discovery and Internal Communication	416
Support for Hostaware Replication	417
Configuring DataGrid Discover URL	419
Multicast (PGM) Cluster Member Discovery	419
Unicast (Well-Known Address) Cluster Member Discovery	420
Configuring DataGrid Listen URL	422
Configuring DataGrid Transport Security	424
Set Up Policy and Token Files	424
Restricting Transport Access	426
Schema/Model Migration with Shared Nothing Persistence	428
CDD Cluster Tab DataGrid Properties Reference	430
Enabling Use of Oracle Coherence as the Cache Provider	433
Configuring Oracle Coherence Cluster Discovery	434
Guidelines for Managing Coherence Clusters	434
Configuring Multicast Cluster Discovery for Coherence Clusters	435
Configuring Well-Known Address Cluster Discovery for Coherence Clusters	435
CDD Cluster Tab Coherence Properties Reference	438
Chapter 25 Load Balancer Configuration	445
Understanding Load Balancing Options	446
Basic Load Balancing	446
Content-aware Load Balancing	446
Content-Aware Load Balancer Configuration	449

Adhoc Load Balancer Configuration	450
Router Configuration	450
Receiver Configuration	451
Receiver Membership Functions in Catalog	451
CDD Load Balancer Tab Properties Reference	452
Chapter 26 Backing Store Configuration	455
Configuring Backing Store Settings and Properties	456
CDD Cluster Tab Backing Store Settings Reference	457
CDD Cluster Tab Backing Store Properties Reference	461
Setting up Shared Nothing Persistence	466
Chapter 27 Berkeley DB Shared All Persistence	469
Berkeley DB Shared All Persistence Overview	470
Configuring the Berkeley DB Shared All Persistence Option	472
Reference to Berkeley DB Shared All Persistence CDD Properties	474
Reference To Berkeley DB (JE) Properties	476
Chapter 28 Domain Objects Configuration	477
Configuring Domain Object Settings	478
Configuring Preloading Options	479
CDD Cluster Tab Domain Objects Default Settings Reference	481
CDD Cluster Tab Domain Object Override Settings Reference	485
Chapter 29 Agent and Processing Unit Configuration	489
Collections, Agent Classes, and Processing Units	490
Collections	490
Agent Classes	490
How Collections and Individual Resources are Used to Configure Agents	490
Log Configurations	491
Configuring Collections of Rules, Rule Functions, and Destinations	492
CDD Collections Tab Input Destination Settings Reference	494
Configuring Agent Classes (All OM Types)	496
CDD Agent Classes Tab Settings Reference	498
CDD Agent Classes Tab Properties Reference	501
Understanding and Configuring Log Configurations	506
Log File Settings	506
Understanding Log Configuration Levels and Syntax	506
Configuring Log Configurations	507

CDD Collections Tab Log Configurations Settings Reference	509
Logging for TIBCO BusinessEvents DataGrid	512
Configuring Processing Units (All OM Types)	513
CDD Processing Units Tab Settings Reference	515
CDD Processing Units Tab Coherence Log Properties Reference	517
CDD Processing Units Tab JMS Server Connection Properties	520
Chapter 30 JDBC Backing Store Setup	523
JDBC Backing Store Setup and Configuration Overview	524
Backing Store Setup	526
Backing Store Configuration	526
Cases That May Need Additional Setup	527
Ontology Identifiers that Exceed the DBMS Maximum Column Length	527
Ontology Identifiers that Use Database Key Words	528
String Properties that Exceed the DBMS Maximum Column Length	528
Excluding Entities from the Backing Store	528
Resources Required for Setting Up the Database	529
Provided Configuration Resources	529
DBMS Software Requirements and Installation	531
SQL Server	532
Oracle Database	534
Configure the CDD for Special Cases (As Needed)	536
Add a JDBC Connection Resource (Now or Later)	537
Configure Backing Store Settings in the CDD (Now or Later)	539
Build the EAR File	542
Initialize the Database and Generate Non-Project Tables	543
Generate the Project-Schema-Specific SQL Scripts	545
Check the Aliases File, Run the Project Schema Script	549
Updating an Existing Backing Store Schema	552
Preparing to Run and Running the Schema Update Utility	552
What the Schema Update Utility Can and Cannot Handle Automatically	553
Handling Entities Deleted from a Write-Behind Backing Store	555
Backing Store Table Reference	556
Primary Tables	556
Secondary Tables	557
Reverse Reference Tables	557
Class-to-Table Mapping	558

Chapter 31 Introduction to WebStudio	559
Introduction to TIBCO BusinessEvents WebStudio	560
Technical User Tasks	560
Business User Task	562
Introduction to Rules Management Server (RMS)	564
WebStudio and RMS User Workflow	566
Chapter 32 RMS and RMS Projects Configuration	567
Adding a Project to RMS	568
Configuring RMS Server Properties	569
Enabling Remote Connection to RMS from WebStudio	571
Enabling SSL Authentication in RMS	572
Enabling One-way SSL Authentication in RMS	572
Enabling Two-way SSL Authentication in RMS	573
Adding Shared JDBC Connection in Runtime	574
RMS Server Configuration Property Reference	575
Chapter 33 Working With Projects in WebStudio	585
Starting and Logging in to RMS	586
Status Check for RMS Connection	586
Checking Out a Project	588
Updating (Synchronizing) a Project	590
Committing Artifacts for Approval	591
Chapter 34 Working with Decision Tables in WebStudio	595
Working With Decision Tables	596
Create/Remove a Decision Table	596
Edit a Decision Table	597
Export a Decision Table	598
Import a Decision Table	599
Understanding Columns and Rows (Rules) in a Decision Table	600
Conditions and Actions	600
Regular and Custom Conditions and Actions	600
Working With Rows and Columns in WebStudio	602
Add Condition and Action Columns	602
Working with Columns	605
Add a Row	606
Setting Rule (Row) Priorities	607
Filter the Rows	608
Remove a Row	608

Supported Operators	609
Using Operators in Tables	610
Analyze a Decision Table	611
Validate a Decision Table	613
Chapter 35 Business Rules	615
Business Rules Overview	616
Working with Business Rules	617
Adding Business Rules	617
Export a Business Rule	620
Deleting Business Rules	620
Business Rule Operators	622
Chapter 36 Working with the Approval Workflow in RMS	625
Working with the Approval Workflow—Overview	626
Approval Status Values	626
Checking a Worklist and Taking Action	628
Chapter 37 Deployable Files Generation and Deployment for WebStudio	631
Generating Deployable Files (EAR and Class Files)	632
Generating the Project EAR or All Project Class Files	633
Generating One Decision Table's Class File	634
Generating One Business Rule's Rule Template Instance File	635
Business Rule Hot Deployment	636
Chapter 38 ActiveMatrix BusinessWorks Integration	637
Overview of Integration with ActiveMatrix BusinessWorks	638
Integration Components	639
ActiveMatrix BusinessWorks Activities Palette	639
TIBCO BusinessEvents Functions	639
Design Considerations	641
Integration Scope	641
Thread Management	641
Configuring a RuleServiceProvider Configuration Resource	642
TIBCO BusinessEvents RuleServiceProvider Configuration Resource Reference	643
Configuration	643
Working With the TIBCO BusinessEvents Activities	645
Receive Event Resource Reference	646
Configuration	646

Misc	647
Output	647
Send Event Resource Reference	648
Configuration	648
Input	648
Wait for Event Resource Reference	649
Configuration	649
Event	650
Input	650
Output	651
Invoking a TIBCO BusinessEvents Rule Function from a Process	652
Specifying Input Arguments	652
Using Synchronous Invocation	652
Using the lockWM Parameter	653
Overriding the Rule Function at Runtime	653
Working With Invoke RuleFunction Activities	654
Invoke RuleFunction Resource Reference	656
Configuration	656
Input	657
Output	657
Working with the BusinessWorks Functions	658
Providing Paths to TIBCO BusinessEvents Project Resources Using Schemas	658
Using invokeProcess()	658
Using startProcess()	660
Using cancelProcess()	662
Using init()	662
Using shutdown()	662
Chapter 39 TIBCO BusinessEvents Performance Profiler	665
Overview of Profiler	666
Changing the Delimiter Character	667
Turning Profiler On and Off	668
To Turn Profiler On and Off Using TIBCO BusinessEvents Monitoring and Management	668
To Turn Profiler On and Off Using Properties	668
To Turn Profiler On and Off Using Functions	671
To Turn Profiler On and Off Using TIBCO Hawk Methods	672
Profiler Reference	674
Chapter 40 Testing and Debugging Projects	679
Overview of Testing and Debugging Projects	680
Debugging Projects	680

Testing Projects	680
Launch Configurations	681
Test Data	681
The Rule Agenda and Variable Views	681
Viewing and Understanding Results	682
Debugger Setup	682
Preparing to Run (Test) or Debug a Project	683
Build an EAR File	683
Create Test Data (as Desired)	683
For Remote Debugging Only, Configure Java Debug Interface (JDI)	683
Adding and Working with Launch (Debug or Run) Configurations	685
Launch Configurations Reference	688
For Testing and Local Debugging	688
For Remote Debugging	689
Creating and Working With Test Data	690
Working with Concept and Event Test Data	690
Working with Rule Data	692
Setting Breakpoints in Rules and Rule Functions	693
Running Debugger	695
Running Tester	696
Asserting Rule Input Data	697
Viewing the Results	699
Understanding Result Data	700
Viewing and Understanding Working Memory Contents	700
Chapter 41 Diagrams	701
Overview of Diagrams	702
Working with Diagrams	704
Configuring Diagram Preferences and Properties	704
Performing Common Tasks	705
Using Diagram Tools	706
Exporting a Diagram to an Image	708
Printing a Diagram	709
Setting Diagram Preferences	711
Project Analyzer and Selected Entity Project Diagrams	712
Advantages of Project Analyzer and Selected Entity Project Diagrams	713
Working with Project Analyzer and Selected Entity Project Diagrams	713
Dependency Diagrams	715
Sequence Diagrams	717
Concept Model Diagrams	718

Event Model Diagrams	719
Diagram Options and Tools Reference	720
Layout Options	722
Chapter 42 TIBCO BusinessEvents Studio Preferences	725
Setting Preferences	726
Preference Sections	726
Decision Table Related Preferences	728
Decision Table Preferences	728
Decision Table Analyzer Preferences	728
Decision Table Appearance Preferences	729
Diagram Preferences	730
Tester Preferences	735
Appendix A Understanding Entity Caches	737
Entity Cache Names Format	737
Caches for Ontology Objects	738
Caches for Internal Entities	738
Appendix B Handling Null Properties	741
Handling Null Concept Property Values	742
Enabling Use of the Nillable Attribute	742
Enabling Null Property Values to Appear When Serializing Concepts to XML	742
Examples of Nillable Attribute and Null Properties Settings	743
Enabling and Setting Special Treatment of Numeric and Boolean Null Values	744
Setting Runtime Properties for Special Treatment of Null Values	745
Property Reference for Null Property Handling	746
Index	749

Tables

Table 1	General Typographical Conventions	xxxvii
Table 2	Syntax Typographical Conventions	xxxviii
Table 3	Global Variable Reference	16
Table 4	Build Enterprise Archive Reference	24
Table 5	TIBCO BusinessEvents Studio Tools Options for Building an EAR File	26
Table 6	TIBCO BusinessEvents Studio Tools Options for Importing an Existing Project	29
Table 7	TIBCO BusinessEvents Studio Tools Options for Working with Project Libraries	32
Table 8	TIBCO BusinessEvents Studio Tools for Generating Class Files	35
Table 9	TIBCO BusinessEvents Studio Tools Options for Generating Encrypted Passwords	37
Table 10	Options for the Command <code>generate</code>	40
Table 11	Options for the Command <code>import</code>	41
Table 12	Options for the Command <code>assert</code>	42
Table 13	Mapping of Coherence Functions to DataGrid Functions	51
Table 14	Refactoring for Move and Rename Operations	54
Table 15	Which Serializers to Use for JMS Message Types	85
Table 16	Variables for Use with <code>DurableSubscriberName</code>	87
Table 17	JMS Message Acknowledgement Modes	89
Table 18	When JMS Messages are Acknowledged	92
Table 19	JMS Header Field Names	95
Table 20	HTTP Channel Advanced Configuration Settings	108
Table 21	Exporting Project Artifacts as WSDL	128
Table 22	Imported WSDL Project Artifacts	136
Table 23	Operators for ActiveSpaces Filters	152
Table 24	Formats for Filter Values	154
Table 25:	Configuring an ActiveSpaces Destination and Event	155
Table 26	Simple Event Payload Element Parameters	170
Table 27	Attributes Used for Each Type of Advisory Event	190
Table 28	Quick Fix Feature Options	271

Table 29	Tips for Working in the Rule Editor	274
Table 30	Common Arguments for VRF Functions	308
Table 31	Function Catalog Elements	318
Table 32	Escape Sequences	326
Table 33	Operators in the TIBCO BusinessEvents Rule Language.	326
Table 34	Attributes	329
Table 35	Concept Properties to XML Datatype Conversions.	340
Table 36	Operator Matrix	341
Table 37	Input tab toolbar buttons	349
Table 38	Icons for schema items	352
Table 39	Additional icons for hints	353
Table 40	Datatype validation	356
Table 41	XPath Formula Builder Reference.	387
Table 42	Formatting characters in date or time strings	391
Table 43	Example date and time format patterns.	393
Table 44	CDD Cluster Tab General Settings	405
Table 45	CDD Cluster Tab Cache OM Settings.	411
Table 46	CDD Cluster Tab TIBCO BusinessEvents DataGrid Properties	430
Table 47	CDD Cluster Tab Coherence Properties	438
Table 48	CDD Load Balancer Tab Properties	452
Table 49	CDD Cluster Tab Backing Store Settings	457
Table 50	CDD Cluster Tab Backing Store Properties	461
Table 51	Berkeley DB Shared All Persistence CDD Configuration Properties.	474
Table 52	Berkeley DB JE Properties	476
Table 53	CDD Cluster Tab Domain Object Default Settings	481
Table 54	CDD Cluster Tab Domain Object Override Settings	485
Table 55	CDD Collections Tab Input Destination Settings	494
Table 56	CDD Agent Classes Tab Inference Agent and Query Agent Settings.	498
Table 57	CDD Agent Classes Tab Inference Agent and Query Agent Properties	501
Table 58	CDD Collections Tab Log Configurations Settings	509
Table 59	CDD Processing Units Tab Settings	515
Table 60	CDD Processing Units Tab Coherence Properties	517

Table 61	CDD Processing Units Tab Properties for Reconnecting to a JMS Server	520
Table 62	Resources Required for JDBC Backing Store Implementation	529
Table 63	Secondary table structure	557
Table 64	RMS Server Configuration Properties	575
Table 65	Supported Operators	609
Table 66	Profiler Configuration Properties	668
Table 67	Profiler Column Heading Reference.	674
Table 68	Types of Diagrams Available for Different Resource Types.	703
Table 69	Context Menu Options for Canvas and Objects.	708
Table 70	TIBCO BusinessEvents Diagram Tools Reference	720
Table 71	Global Diagram Preferences	730
Table 72	Specific Diagram Preferences (Alphabetical).	732
Table 73	Reference to Tester Preferences.	735
Table 74	Reference to Tester Appearance Preferences.	735
Table 75	Internal Entity Caches	738
Table 76	Properties for Null Property Handling.	746

Preface

TIBCO BusinessEvents® allows you to abstract and correlate meaningful business information from the events and data flowing through your information systems, and take appropriate actions using business rules. By detecting patterns within the real-time flow of events, TIBCO BusinessEvents can help you to detect and understand unusual activities as well as recognize trends, problems, and opportunities. TIBCO BusinessEvents publishes this business-critical information in real time to your critical enterprise systems or dashboards. With TIBCO BusinessEvents you can predict the needs of your customers, make faster decisions, and take faster action.

Topics

- [Changes from the Previous Release of this Guide, page xxvi](#)
- [TIBCO BusinessEvents Express, page xxxi](#)
- [Related Documentation, page xxxii](#)
- [Typographical Conventions, page xxxvii](#)
- [Connecting with TIBCO Resources, page xl](#)

Changes from the Previous Release of this Guide

This section itemizes the major changes from the previous release of this guide.

Release 5.1.2 - Document Updated: July 2014

The current release is set everywhere in this document to 5.1.2.

Release 5.1.2

TIBCO BusinessEvents Studio Tools Utility

- Updated buildEar operation syntax for `-pl`, `-lc` and `-cp` options. See [Building an Enterprise Archive \(EAR File\) at the Command Line on page 26](#).
- Updated generateClass operation syntax for `-pl` and `-lc` options. See [Generating All Project Class Files at the Command Line on page 35](#).

HTTP Channels

- Removed sessionConcept from sample code for action rule function. See [Working with HTTP Requests on page 99](#).
- Added the section [Setting Up Fault Tolerance for the HTTP Channel on page 106](#).
- Added note for for configuring SSL over HTTP when the Identity Resource is of the type Certificate/KeyURL. See [Using HTTP Functions to Configure HTTP Request Messages on page 116](#).

TIBCO BusinessEvents WebStudio

- Added step to store custom function for your WebStudio project. See [Adding a Project to RMS on page 568](#).
- Added the section [Enabling SSL Authentication in RMS on page 572](#).
- Added the section [Adding Shared JDBC Connection in Runtime on page 574](#).
- Updated the RMS.cdd properties reference. See [RMS Server Configuration Property Reference on page 575](#).
- Added the section [Business Rule Hot Deployment on page 636](#).

Miscellaneous

- Section [Setting up Shared Nothing Persistence on page 466](#) as been updated to list the reserved words for Shared Nothing: extid, id, closure, next, tti, and fired.

These words cannot be used as a name of a property in an event or concept.

Release 5.1

Documentation Restructuring

Chapters relating to setting up a cluster and a backing store have been moved from the *TIBCO BusinessEvents Administration* guide to *TIBCO BusinessEvents Developer's Guide* (this guide). Now the *TIBCO BusinessEvents Administration* guide contains only deploytime and runtime administration topics.

Studio-tools Changes

Topics relating to studio-tools command-line utilities used in TIBCO BusinessEvents Decision Manager are now shown only in *TIBCO BusinessEvents Decision Manager User's Guide*. Command-line syntax has also changed in this release. Documentation for tools in the studio-tools suite is provided in relevant sections of the guide.

Default for ParallelOps Changed

In release 5.0, parallel operations was enabled by default when cache aside database write strategy was used. Now parallel operations is enabled by default only when both cache aside *and* concurrent RTC features are used. See [Table 56, CDD Agent Classes Tab Inference Agent and Query Agent Settings, on page 498](#).

Hawk Channel

The Hawk channel connects TIBCO BusinessEvents to a TIBCO Hawk domain and can be configured to use either TIBCO Rendezvous or TIBCO Enterprise Message Service transports. The Hawk channel enables TIBCO BusinessEvents to receive events from the Hawk monitor and transform them into events. Information about how to configure and work with the channel has been added in [Chapter 7, HAWK Channel, on page 141](#).

Connection to TIBCO Hawk domain through a specific transport is configured using the shared resource Hawk connection. See [Hawk Connection on page 223](#).

ActiveSpaces Channel

The ActiveSpaces channel connects TIBCO BusinessEvents to TIBCO ActiveSpaces metaspace. This enables TIBCO BusinessEvents to monitor the activities on the TIBCO ActiveSpaces metaspace and receive events from TIBCO ActiveSpaces and convert them into events in TIBCO BusinessEvents. Information about how to configure and work with the channel has been added in [Chapter 8, ActiveSpaces Channel, on page 148](#).

Connection to a space in TIBCO ActiveSpaces is configured using the shared resource ActiveSpaces connection. See [ActiveSpaces Connection on page 220](#).

HTTP Channel Changes

The BUILT-IN server option is no longer required and has been removed. Tomcat server version 7 is now supported, and it provides functionality previously provided by the BUILT-IN server type.

Backing Store Persistence Options

The backing store can now be configured with one of the following persistence options: None, Shared All, or Shared Nothing. Chapters related to backing store configuration have been updated with information about the new persistence options.

Load Balancer Configuration for Agents

Agents can be configured to work cooperatively as routers and receivers to ensure that related messages arriving from queue sources are handled by the same agent, so that related information is locally available. [Chapter 25, Load Balancer Configuration](#) explains the load balancing options and provides the steps to configure from the CDD file.

Generate JDBC Deployment Scripts From TIBCO BusinessEvents Studio

TIBCO BusinessEvents Studio provides an alternative approach to generate the project-schema-specific SQL scripts using the **JDBC Deployment** wizard. The wizard also allows you to generate SQL scripts to migrate existing backing store database. See [Generating Scripts Using the JDBC Deployment Wizard on page 545](#) for details.

TIBCO BusinessEvents WebStudio and Rules Management Server

TIBCO BusinessEvents WebStudio is an online tool to define business rules and Rule Management Server (RMS) is its server component. Earlier RMS component was part of TIBCO BusinessEvents Decision Manager add-on only but now its part of the TIBCO BusinessEvents Standard. Following chapters are added in the *TIBCO BusinessEvents WebStudio User's Guide* for this release:

- [Chapter 17, Rule Templates, page 279](#)
- [Chapter 31, Introduction to WebStudio, page 559](#)
- [Chapter 32, RMS and RMS Projects Configuration, page 567](#)
- [Chapter 33, Working With Projects in WebStudio, page 585](#)
- [Chapter 34, Working with Decision Tables in WebStudio, page 595](#)
- [Chapter 35, Business Rules, page 615](#)
- [Chapter 36, Working with the Approval Workflow in RMS, page 625](#)
- [Chapter 37, Deployable Files Generation and Deployment for WebStudio, page 631](#)

Miscellaneous

- The Berkeley DB persistence option has been removed from the product. All mention of Berkeley DB persistence has been removed from documentation.
- You can now use global variables in a project library and order project libraries to specify an override order for global variables that are listed more than one time. See [Setting and Overriding Global Variables from a Project Library on page 14](#).
- Better information on post RTC to backing store tuning properties.
- Better information on backing store commit property.
- Domain model entries are case sensitive. For example, m and M are recognized as different entries. This is now noted.
- New information is added regarding JMS message acknowledgement. See [Using CLIENT_ACKNOWLEDGE Mode with Websphere MQ and Cache-Aside on page 91](#).

- Terminology for certain threading options has been rationalized. Changes are shown in bold in the following table:

Old Name	New Name
<ul style="list-style-type: none">• Shared Queue<ul style="list-style-type: none">— Size— Workers	<ul style="list-style-type: none">• Shared Queue<ul style="list-style-type: none">— Queue Size— Thread Count
<ul style="list-style-type: none">• Workers<ul style="list-style-type: none">— Queue Size— Thread Count	<ul style="list-style-type: none">• Destination Queue<ul style="list-style-type: none">— Queue Size— Thread Count

- A new property is added to specify how a concept is stored in the TIBCO BusinessEvents DataGrid cache. See [CDD Cluster Tab DataGrid Properties Reference on page 430](#).

TIBCO BusinessEvents Express

The TIBCO BusinessEvents Express edition provides more limited functionality than the TIBCO BusinessEvents Standard Edition. Some content in this documentation is not relevant to users of TIBCO BusinessEvents Express. Such content includes but is not limited to any chapters and major sections that contain a note indicating that the content does not apply to TIBCO BusinessEvents Express.

Minor references to unsupported features may not be called out in the text. Use the following general guidelines to understand what is and is not supported in these cases:

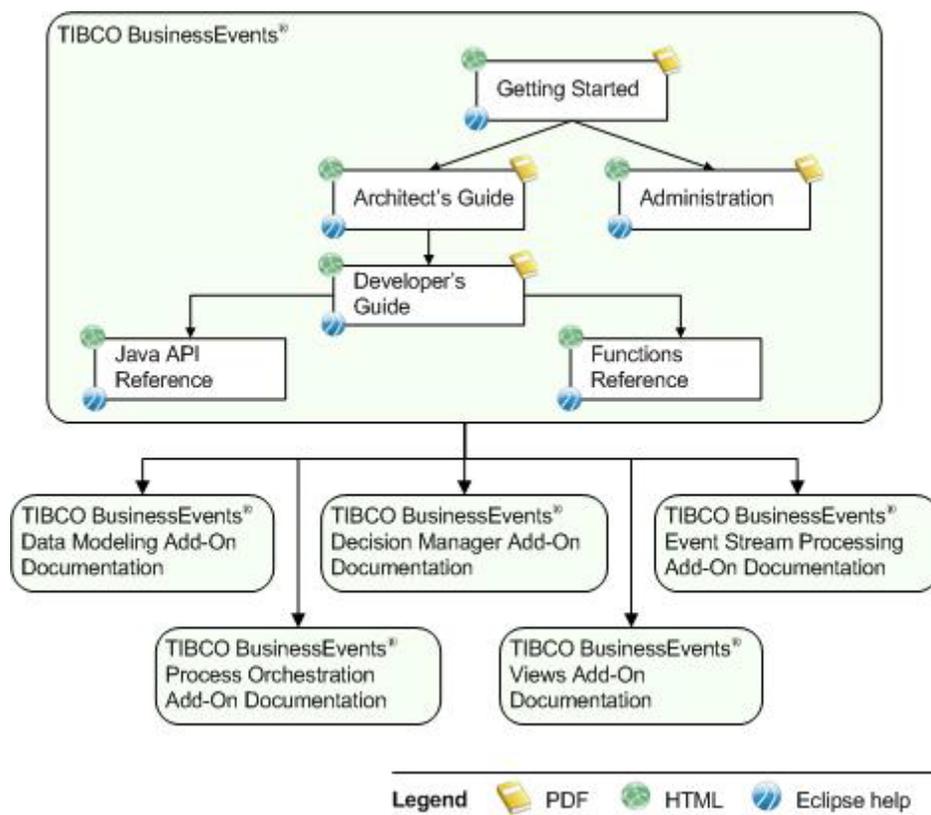
- Only In Memory object management (OM) is supported. Therefore all functionality that requires Cache OM, such as use of a backing store, is not available. Berkeley DB OM is also not supported with the TIBCO BusinessEvents Express edition.
- Only the TIBCO BusinessEvents Decision Manager add-on is supported with the TIBCO BusinessEvents Express edition in this release. Other add-on products are not supported.

Related Documentation

This section lists documentation resources you may find useful.

TIBCO BusinessEvents and Add-On Product Documentation

The following diagram shows the main documents in the TIBCO BusinessEvents documentation set, and the documentation sets for the optional add-on products.



Each set also contains an installation guide, release notes, and a readme file.

TIBCO BusinessEvents Documentation

TIBCO BusinessEvents Studio, the design-time UI, is supported on Windows and Linux. The documentation set for TIBCO BusinessEvents is as follows.

- *TIBCO BusinessEvents Installation*: Read this manual for instructions on site preparation, installation, upgrading from an earlier release, and project migration.
- *TIBCO BusinessEvents Getting Started*: After the product is installed, use this manual to learn the basics of TIBCO BusinessEvents: project design, cache OM, and backing store. This guide explains the main ideas so you gain understanding as well as practical knowledge.
- *TIBCO BusinessEvents Architect's Guide*: If you are architecting an application using TIBCO BusinessEvents, read this guide for overview and detailed technical information to guide your work.
- *TIBCO BusinessEvents Developer's Guide*: Use this guide when you implement a project design in TIBCO BusinessEvents Studio. It covers topics such as project-level tasks, resource-level tasks, debugging, and integration with TIBCO ActiveMatrix BusinessWorks. It also explains how to configure the CDD file for different object management options, and set up a backing store.
- *TIBCO BusinessEvents Administration*: This book explains how to configure, deploy, monitor, and manage a TIBCO BusinessEvents application and the data it generates using TIBCO BusinessEvents Monitoring and Management component, TIBCO Administrator, or at the command line. It includes authentication and authorization topics.
- Online References:
 - *TIBCO BusinessEvents Java API Reference*: This online reference is available from the HTML documentation interface. It provides the Javadoc-based documentation for the TIBCO BusinessEvents API.
 - *TIBCO BusinessEvents Functions Reference*: This reference is available from the HTML documentation interface. It provides a listing of all functions provided with TIBCO BusinessEvents, showing the same details as the tooltips available in TIBCO BusinessEvents Studio.
- *TIBCO BusinessEvents Release Notes*: Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

TIBCO BusinessEvents Event Stream Processing

This TIBCO BusinessEvents add-on is available separately, and includes the TIBCO BusinessEvents Query Language features and the Pattern Matcher Service.

- *TIBCO BusinessEvents Event Stream Processing Installation*: Read this brief manual for installation instructions. A compatible version of TIBCO BusinessEvents must be installed before you install any add-on.

- *TIBCO BusinessEvents Event Stream Processing Query Developer's Guide*: This manual explains how to use the object query language to query various aspects of the running system. For details on configuring and deploying query agents, see *TIBCO BusinessEvents Developer's Guide*.
- *TIBCO BusinessEvents Event Stream Processing Pattern Matcher Developer's Guide*: This manual explains how to use the pattern matcher language and engine to correlate event patterns in a running system.
- *TIBCO BusinessEvents Event Stream Processing Release Notes*: Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

TIBCO BusinessEvents Decision Manager

This TIBCO BusinessEvents add-on is available separately. It incorporates the Decision Manager decision modeling business user interface (supported on Windows and Linux), and the Rules Management Server (supported on all platforms supported by TIBCO BusinessEvents).

- *TIBCO BusinessEvents Decision Manager Installation*: Read this brief manual for installation instructions. A compatible version of TIBCO BusinessEvents must be installed before you install any add-on.
- *TIBCO BusinessEvents Decision Manager User's Guide*: This manual explains how business users can use decision tables and other decision artifacts to create business rules. It also covers configuration and administration of Rules Management Server, which is used for authentication, authorization, and approval processes.
- *TIBCO BusinessEvents Decision Manager Release Notes*: Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

TIBCO BusinessEvents Data Modeling

This TIBCO BusinessEvents add-on is available separately. It contains state models and database concept features.

- *TIBCO BusinessEvents Data Modeling Installation*: Read this brief manual for installation instructions. A compatible version of TIBCO BusinessEvents must be installed before you install any add-on.
- *TIBCO BusinessEvents Data Modeling Developer's Guide*: This manual explains data modeling add-on features for TIBCO BusinessEvents. The database concepts feature enables you to model TIBCO BusinessEvents concepts on Database tables. The state modeler feature enables you to create state machines.

- *TIBCO BusinessEvents Data Modeling Release Notes*: Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

TIBCO BusinessEvents Process Orchestration

This TIBCO BusinessEvents add-on is available separately. It provides CEP functionality within the context of a BPM process, enabling you segregate different CEP rule sets within the flow of a BPM process.

- *TIBCO BusinessEvents Process Orchestration Installation*: Read this manual for instructions on site preparation and installation. A compatible version of TIBCO BusinessEvents must be installed before you install any add-on.
- *TIBCO BusinessEvents Process Orchestration Developer's Guide*: This guide explains how configure and deploy business processes whose actions are carried out using TIBCO BusinessEvents project resources.
- *TIBCO BusinessEvents Process Orchestration Release Notes*: Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

TIBCO BusinessEvents Views

This TIBCO BusinessEvents add-on is available separately. It includes graphical dashboard components for run-time event monitoring.

- *TIBCO BusinessEvents Views Installation*: Read this manual for instructions on site preparation and installation. A compatible version of TIBCO BusinessEvents must be installed before you install any add-on.
- *TIBCO BusinessEvents Views Getting Started*: After the product is installed, use this manual to learn how to use TIBCO BusinessEvents Views to create and run a dashboard using a step-by-step tutorial.
- *TIBCO BusinessEvents Views Developer's Guide*: This guide explains how to use TIBCO BusinessEvents Views to create meaningful metrics that are presented to business users in real-time for proactive decision making.
- *TIBCO BusinessEvents Views User's Guide*: This book explains how to monitor metrics in TIBCO BusinessEvents TIBCO BusinessEvents Views and how to represent the business processes graphically.
- *TIBCO BusinessEvents Views Release Notes*: Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

Accessing TIBCO BusinessEvents Functions Reference Documentation

Reference documentation for functions, including those used in add-ons, is available in the HTML documentation interface for the TIBCO BusinessEvents documentation set, and as tooltips in TIBCO BusinessEvents Studio. To use the HTML-based functions reference from the file system do the following:

1. Browse to *BE_HOME/doc/standard/html* and click **index.htm**. The HTML documentation interface appears.
2. In the left panel, browse to Online References and in the right panel choose TIBCO BusinessEvents Functions Reference. The reference opens in a new tab.
3. Click the navigation links to browse to the functions as desired.

Other TIBCO Product Documentation

You may find it useful to refer to the documentation for the following TIBCO products:

- TIBCO ActiveSpaces[®]
- TIBCO Hawk[®]
- TIBCO Rendezvous[®]
- TIBCO Enterprise Message Service[™]
- TIBCO ActiveMatrix BusinessWorks[™]

Typographical Conventions

The following typographical conventions are used in this manual.

Table 1 General Typographical Conventions

Convention	Use
<i>ENV_NAME</i> <i>TIBCO_HOME</i> <i>BE_HOME</i>	<p>TIBCO products are installed into an installation environment. A product installed into an installation environment does not access components in other installation environments. Incompatible products and multiple instances of the same product must be installed into different installation environments.</p> <p>An installation environment consists of the following properties:</p> <ul style="list-style-type: none"> • Name Identifies the installation environment. This name is referenced in documentation as <i>ENV_NAME</i>. On Microsoft Windows, the name is appended to the name of Windows services created by the installer and is a component of the path to the product shortcut in the Windows Start > All Programs menu. • Path The folder into which the product is installed. This folder is referenced in documentation as <i>TIBCO_HOME</i>. <p>TIBCO BusinessEvents installs into a directory within a <i>TIBCO_HOME</i>. This directory is referenced in documentation as <i>BE_HOME</i>. The default value of <i>BE_HOME</i> depends on the operating system. For example on Windows systems, the default value is C:\tibco\be\5.1.</p>
code font	<p>Code font identifies commands, code examples, filenames, pathnames, and output displayed in a command window. For example:</p> <p>Use <code>MyCommand</code> to start the foo process.</p>
bold code font	<p>Bold code font is used in the following ways:</p> <ul style="list-style-type: none"> • In procedures, to indicate what a user types. For example: Type admin. • In large code samples, to indicate the parts of the sample that are of particular interest. • In command syntax, to indicate the default parameter for a command. For example, if no parameter is specified, <code>MyCommand</code> is enabled: <code>MyCommand [enable disable]</code>

Table 1 General Typographical Conventions (Cont'd)

Convention	Use
<i>italic font</i>	<p>Italic font is used in the following ways:</p> <ul style="list-style-type: none"> To indicate a document title. For example: See <i>TIBCO ActiveMatrixBusinessWorks Concepts</i>. To introduce new terms. For example: A portal page may contain several <i>portlets</i>. Portlets are mini-applications that run in a portal. To indicate a variable in a command or code syntax that you must replace. For example: <code>MyCommand <i>PathName</i></code>
Key combinations	<p>Key name separated by a plus sign indicate keys pressed simultaneously. For example: <code>Ctrl+C</code>.</p> <p>Key names separated by a comma and space indicate keys pressed one after the other. For example: <code>Esc, Ctrl+Q</code>.</p>
	The note icon indicates information that is of special interest or importance, for example, an additional action required only in certain circumstances.
	The tip icon indicates an idea that could be useful, for example, a way to apply the information provided in the current section to achieve a specific result.
	The warning icon indicates the potential for a damaging situation, for example, data loss or corruption if certain steps are taken or not taken.

Table 2 Syntax Typographical Conventions

Convention	Use
[]	<p>An optional item in a command or code syntax.</p> <p>For example:</p> <pre>MyCommand [optional_parameter] required_parameter</pre>
	<p>A logical OR that separates multiple items of which only one may be chosen.</p> <p>For example, you can select only one of the following parameters:</p> <pre>MyCommand param1 param2 param3</pre>

Table 2 *Syntax Typographical Conventions*

Convention	Use
{ }	<p>A logical group of items in a command. Other syntax notations may appear within each logical group.</p> <p>For example, the following command requires two parameters, which can be either the pair <code>param1</code> and <code>param2</code>, or the pair <code>param3</code> and <code>param4</code>.</p> <pre>MyCommand {param1 param2} {param3 param4}</pre> <p>In the next example, the command requires two parameters. The first parameter can be either <code>param1</code> or <code>param2</code> and the second can be either <code>param3</code> or <code>param4</code>:</p> <pre>MyCommand {param1 param2} {param3 param4}</pre> <p>In the next example, the command can accept either two or three parameters. The first parameter must be <code>param1</code>. You can optionally include <code>param2</code> as the second parameter. And the last parameter is either <code>param3</code> or <code>param4</code>.</p> <pre>MyCommand param1 [param2] {param3 param4}</pre>

Connecting with TIBCO Resources

This section provides links to helpful TIBCO resources.

How to Join TIBCOCommunity

TIBCOCommunity is an online destination for TIBCO customers, partners, and resident experts, a place to share and access the collective experience of the TIBCO community. TIBCOCommunity offers forums, blogs, and access to a variety of resources. To register, go to <http://www.tibcommunity.com>.

How to Access TIBCO Documentation

You can access TIBCO documentation here:

<http://docs.tibco.com>

How to Contact TIBCO Support

For comments or problems with this manual or the software it addresses, contact TIBCO Support as follows:

- For an overview of TIBCO Support, and information about getting started with TIBCO Support, visit this site:

<http://www.tibco.com/services/support>

- If you already have a valid maintenance or support contract, visit this site:

<https://support.tibco.com>

Entry to this site requires a user name and password. If you do not have a user name, you can request one.

Chapter 1 **Project Tasks**

This chapter covers various project-level tasks for TIBCO BusinessEvents Studio, and presents a short introduction to project development.

For information about analyzing or viewing projects, see [Chapter 41, Diagrams, on page 701](#).

Topics

- [Project Development Overview, page 2](#)
- [Creating a Project, page 3](#)
- [Finding a Project Element, page 4](#)
- [Exporting \(Generating\) Concept and Event Schema \(XSD\) Files, page 5](#)
- [Validating a Project or Project Resource, page 7](#)
- [Working with External Library and Custom Function Paths, page 8](#)
- [Working with Project Libraries, page 10](#)
- [Working with Global Variables, page 14](#)
- [Storing Trusted Certificates Outside of Your Project, page 20](#)
- [Building an Enterprise Archive \(EAR File\), page 22](#)

Project Development Overview

TIBCO BusinessEvents Developer's Guide explains how to build a TIBCO BusinessEvents project using the TIBCO BusinessEvents core project resources such as channels, events, concepts, rules, and so on. This brief introductory chapter places this guide in the context of other documentation so you can understand its role in the overall documentation set.

TIBCO BusinessEvents Studio is an Eclipse-based UI used to design, build, and maintain TIBCO BusinessEvents projects. It is integrated into the standard Eclipse menus where appropriate, and works with many established Eclipse UI methodologies and plug-ins.

TIBCO BusinessEvents WebStudio is an online component which allows business users to create/manage business rules in a web browser. In it the user defines an executable rule (business rule) based on the rule template and the rule template view defined by the developer in TIBCO BusinessEvents Studio.

TIBCO BusinessEvents WebStudio is explained in the document *TIBCO BusinessEvents® WebStudio User's Guide*.

A TIBCO BusinessEvents Studio project contains resources used to build, test, and view a design-time TIBCO BusinessEvents project.

Before you begin to use *TIBCO BusinessEvents Developer's Guide*, gain a basic familiarity with the product by completing the tutorials in *TIBCO BusinessEvents Getting Started*.

TIBCO BusinessEvents Developer's Guide provides practical details on using TIBCO BusinessEvents Studio to build a project and configure its cluster deployment descriptor.

For in-depth explanations about designing a TIBCO BusinessEvents project and other topics, read *TIBCO BusinessEvents Architect's Guide*. References to relevant topics in that guide are provided in *TIBCO BusinessEvents Developer's Guide*.

When it is time to configure your application for deployment, deploy, and manage it, refer to *TIBCO BusinessEvents Administration*.

Creating a Project

This section explains the basic procedure for creating a new project. See *TIBCO BusinessEvents Getting Started* for tutorials on building and populating a basic TIBCO BusinessEvents project.



To import a TIBCO BusinessEvents Studio project into your workspace, select File > Import > General > Existing projects into Workspace. Do this for current release projects only.

To import a project from a prior release, see Chapter 4, Migrating Projects from Earlier Versions in *TIBCO BusinessEvents Installation*.

To Create a Project

1. Start TIBCO BusinessEvents Studio. In Windows, click Start > All Programs > TIBCO BusinessEvents Studio.
2. If prompted, select or create the Eclipse workspace directory where your project files will be stored. (If you check the option to use this workspace as a default, you are not prompted again.)
3. You see the TIBCO BusinessEvents Studio UI. From the File menu select **New > Project > TIBCO BusinessEvents > Studio Project** and click **Next**.
4. In the New Studio Project field, enter a project name.

Names follow Java variable naming restrictions. Do not use any reserved words. See [Identifier Naming Requirements on page 323](#).

5. Accept the default location or uncheck the Use default location checkbox and specify the desired location.
6. Click **Finish**.

A new folder tree appears in BusinessEvents Studio Explorer with a default set of folders you can use as desired. (The `defaultVars` folder, however, is required.)



Source Control If the project is under source control using Perforce, editing a resource automatically checks out the resource and makes it writable.

If the file is not under Perforce, but is read only, you are prompted to make the file writable when a change is made.

Finding a Project Element

You can find a project element in various ways and then take some action relating to that element such as opening its editor, or creating a dependency diagram. Depending on the size and complexity of the project, you may find some of the below methods of finding elements more convenient than others.

Note that an element can be a part of a resource (such as a concept property):

- Expand the project tree in BusinessEvents Studio Explorer and double-click the element name.
- Select **Navigate > Open Studio Element** and select an element from the alphabetical list of elements. You can search using wild cards.
- Select the **Resource** perspective, then select **Navigate > Open Resource** (or press **Ctrl+Shift+R**). A dialog allowing you to search through all resources in your workspace appears. To display the whole list, enter ******. You can also use the asterisk (*****) as a wild card. For example to display all concept resources, enter ***.concept**.
- Open the **Project View** diagram and navigate to the element in the diagram. You can search using filters. Double click the element on the diagram to open it. See [Chapter 41, Diagrams, on page 701](#).

Exporting (Generating) Concept and Event Schema (XSD) Files

The Generate Schema utility lets you export concepts and events to XML Schema Definition (XSD) files, one per entity, in a specified location. You can generate schema for all concepts, all events, or both. (You cannot generate schema for one selected entity type.) The files use the same folder structure as the project from which they are exported. In addition, `_BaseConcept.xsd` and `_BaseEvent.xsd` are generated in the root of the selected directory.

XML schemas are used for interoperability between TIBCO BusinessEvents and third party tools or SOA platforms that use well-defined XML for message communication, transformation, and validation.

In the XSD files, concepts are represented as follows:

- Each concept is exported to its own complex type using its own namespace.
- Referenced concepts have a corresponding `ref` attribute in the parent's complex type.
- Contained concepts have a corresponding `type` attribute in the parent's complex type.

To Generate an XML Schema (XSD File)

1. Select the project whose schema you want to generate.
2. From the File menu select **Export**. In the Export wizard, expand **TIBCO BusinessEvents** and select **Generate Schema**.

You can also reach this utility from the option (right-click) menu anywhere in BusinessEvents Studio Explorer.

3. In the Schemas Folder field, browse to and select the folder where you want to put the schema files.
4. Select the **Override TIBCO BusinessEvents Namespace** checkbox to specify a different namespace.

If you do not select a different namespace, an informational message displays. Click **Yes** to continue or **No** to return to the dialog and provide a namespace.

Provide a different namespace to avoid conflicts with the source concepts and events. If you do not provide a namespace, the default TIBCO BusinessEvents namespace is used. (TIBCO BusinessEvents events and concepts have a hidden schema.) If the source entities and generated schema files are in the same folders, use of the default TIBCO BusinessEvents namespace results in a namespace conflict. In this case, you must provide a namespace.

5. In the Select Resources panel, select Concepts or Events or both. Schemas for all concepts or all events in the project (or both) are generated accordingly.
6. Click **Finish**.

The XSD files for the selected resources (all concepts, all events, or both) are generated in subdirectories of the selected directory. Subdirectory names match the project folders. The `_BaseConcept.xsd` and `_BaseEvent.xsd` files are generated in the root of the selected directory.

Validating a Project or Project Resource

When you save a resource, TIBCO BusinessEvents performs validation checks to ensure that all resource requirements are met. For example, it checks that required fields are completed, names are valid, syntax in rules is correct and no unknown functions are called.

You can also perform validation explicitly for an entire project or selected project resource

To Validate a Project, Project Folder, or Project Resource

In BusinessEvents Studio Explorer, do one of the following:

- Right click a project name, folder name, or a project resource name, and select **Validate Project**.
- Select a project name, folder name, or a project resource name and select **Project > Validate Project**.

A pop-up window displays the message, “Validation was successful” or summary information about any problems. Details about problems display in the Problems view.

To Fix Validation Errors

Many validation issues can be fixed using the Quick Fix feature. In the Problems view, right click on a problem and select Quick Fix.

See also [Using the Quick Fix Feature in the Rule Editor on page 271](#) for a related use of this feature.

Project Validation and Project Analysis

In addition to validating for internal consistency, you can run a project analyzer feature. It helps you to understand your project better, and find ways to improve it. For example, it identifies rule functions that are not used in any rule, and rules that will never be triggered at runtime. See [Project Analyzer and Selected Entity Project Diagrams on page 712](#).

Working with External Library and Custom Function Paths

When you work with external libraries or custom functions in your project at design time, and when you run, test, or debug such projects in TIBCO BusinessEvents Studio, you must ensure that the engine can find all libraries (including dependencies on third-party libraries) and custom functions, if the project requires any.

If a project uses Rendezvous channels, for example, Java allows the user to set a property, `java.library.path`, to ensure that the environment path contains a reference to the directory containing Rendezvous DLLs (or on UNIX the SO files).

Similarly TIBCO BusinessEvents captures the native library path along with the build path information to pass to engines running inside TIBCO BusinessEvents Studio. You can enter this information as described below, or when you are configuring a run configuration or debug configuration.



Test Connection Feature To make the Test Connection feature work for JMS Connection and JDBC Shared Connection shared resources, see [Enabling the Test Connection Feature on page 219](#).

To Add an External Library or Custom Function Path



This functionality is also available in the Debug Configurations > ClassPath tab, and the Run Configurations > ClassPath tab, for your convenience. See [Chapter 40, Testing and Debugging Projects, on page 679](#).

1. Open the project in TIBCO BusinessEvents Studio.
2. In BusinessEvents Studio Explorer, right-click the project and click **Properties** (or press Alt-Enter). You see the properties dialog for the project.
3. In the left panel, select **Build Path** and then select one of the following:
 - The **Custom Functions** tab, to add custom functions paths
 - The **Java Libraries** tab, to add third-party libraries paths
4. Click **Add Library** and browse to and select the JAR file.

5. If a library or custom function depends on additional third-party Java archives, do the following:
 - a. Expand the library or custom function entry and double-click the **Native library location** entry.
 - b. At the Native Library Folder Configuration dialog, select **External Folder** or **Workspace** as appropriate, and browse to and select the folder containing the native libraries.
 - c. Click **OK** to dismiss this dialog.
6. Click **OK**. One of the following occurs:
 - If a custom function depends on third-party Java archives and you did not do [step 5](#), a dialog displays when you click OK. It lists the name of the class or classes that could not be loaded. To resolve the problem, do [step 5](#).
 - You see the message: "It is recommended that you rebuild your project after changing the build path. Would you like to rebuild now?" Click Yes to validate and fix any references already in code to the JARs you have added.
7. Save the resource.



Deployment Action Required To make custom functions or external libraries available at runtime, do the following on all machines where TIBCO BusinessEvents is installed:

- Either copy the JAR or JARs to the `lib/ext/tpcl` directory, or other directory in the classpath; or update the classpath in the TRA file to point to the location of the JAR or JARs.
- If a JAR has dependencies on native libraries, edit `BE_HOME/bin/be-engine.tra` and as needed, update `PATH`, `LD_LIBRARY_PATH`, `SHLIB_PATH`, and `LIBPATH` as needed, depending on the operating system.

See [Updating Classpath, Environment Variables, Path, and Copying JAR Files in TIBCO BusinessEvents Administration](#).

To Remove an External Library or Custom Functions JAR Path

Follow the instructions in [To Add an External Library or Custom Function Path](#), but instead of clicking Add Library, select the JAR file you want to remove and click Remove Library. The entry for any native libraries it depends on is also removed.

Working with Project Libraries

This section explains how to work with project libraries in TIBCO BusinessEvents Studio. You can also work with project libraries at the command line (see [Working With Project Libraries at the Command Line on page 31](#) for details).

Project libraries (design-time libraries) are archives that enable you to create project resources once, and share them with other projects.

Project libraries can contain any resources from a TIBCO BusinessEvents Studio project. For example, a project library might contain some concepts that are standard across projects (such as Person), so that multiple projects do not need to redefine these concepts.

Project libraries can include (or, as desired, can consist only of) global variables.

Project libraries have the file extension `.projlib`. TIBCO BusinessEvents Studio includes the following features to allow the refactoring of project resources into project libraries.



Avoid Conflicting Elements Ensure that the project library does not use elements already in use in the local project. If elements in an imported project library conflict with the local project elements, those in the local project have priority. The project library element is ignored in this case for build purposes and is not included in the EAR file.

If multiple project libraries have the same element, then the first one on the build path wins, and the other ones are ignored.

Creating (Exporting) a Project Library

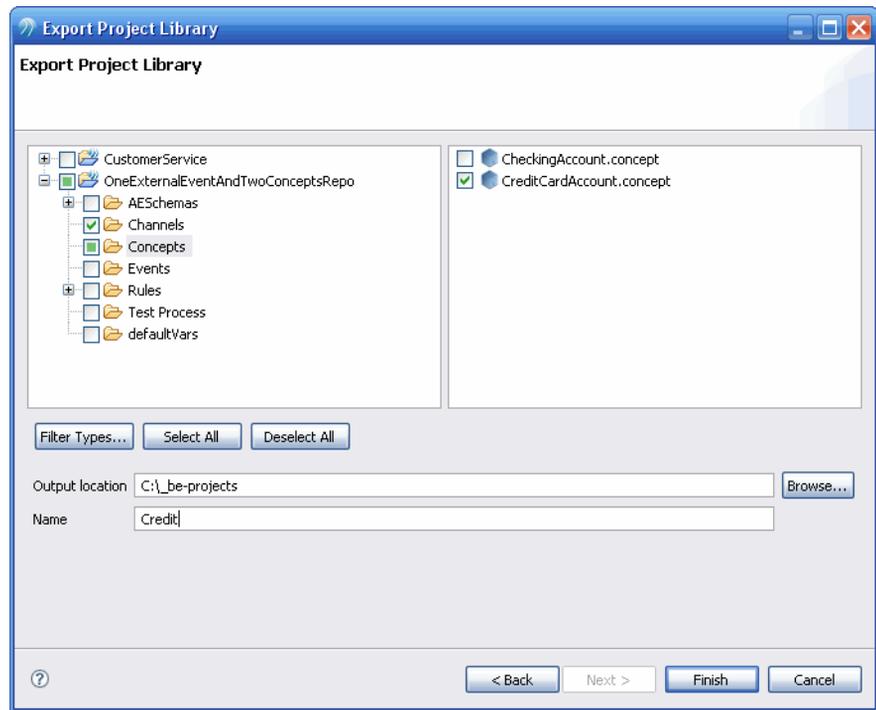
To create a project library you export selected resources to an archive.

To Create a Project Library

1. Optionally, select a project or any of its resources in TIBCO BusinessEvents Studio.
2. Do one of the following:
 - From the File menu select **Export**.
 - Right-click anywhere in BusinessEvents Studio Explorer and select **Export**.

In the Select dialog, select **TIBCO BusinessEvents > Project Library** and click **Next**. You see the Export Studio Project Library dialog.

- In the Export Studio Project Library dialog, select the resources you want to export. (If you selected resources in BusinessEvents Studio Explorer they are selected automatically.)



The left panel displays folders. The right panel displays resources. Select what to export in any of the following ways:

- Select a folder to select all resources within it.
 - Highlight a folder to display its resources in the right panel. Select individual resources as desired.
 - To reduce the set of selected resources, first select resources then click the Filter Types button. Select one or more resource types, then click **OK**. The Export Studio Project Library dialog selection is now reduced to only resources of the selected type or types.
- Enter an output location and give the project library a name, then click **Finish**.



Remember to include the library you just exported, to avoid validation issues.

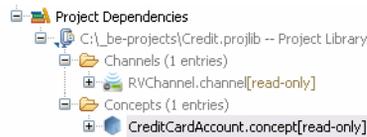
Adding (and Removing) Project Libraries in a Project

Adding a project library does not physically import the files. Instead a pointer to the files is maintained in the project and the resources appear in BusinessEvents Studio Explorer as if they are in the project. The physical location can be in the project folders or external to them, as long as they are available to the project at design time.

You can add a project library to a project in either of these two ways:

- By including a project library
- By importing a project library

After you add the project library, it appears at the root of the project tree as a Project Dependencies node.



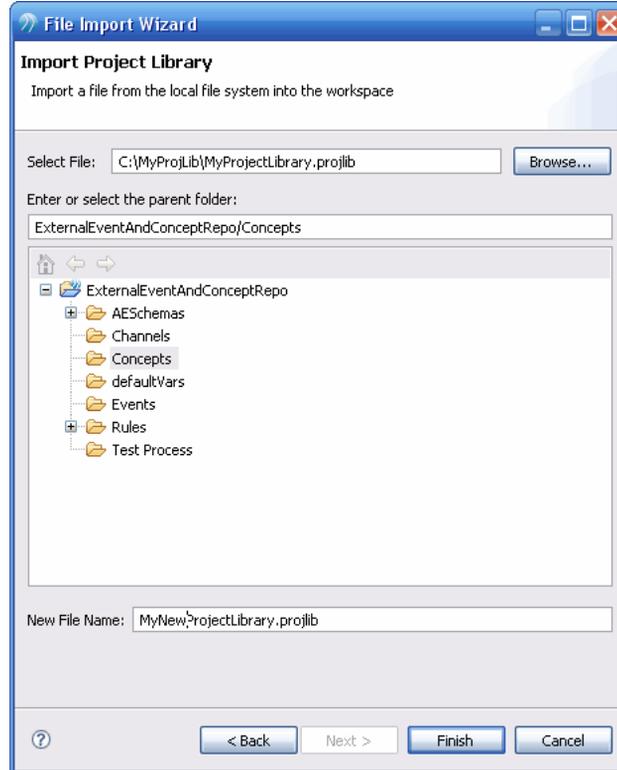
To Include a Project Library (Add it to the Build Path)

1. Open the project in TIBCO BusinessEvents Studio.
2. In BusinessEvents Studio Explorer, right-click the project name and click **Properties** (or press Alt-Enter, or select Project > Properties). You see the properties dialog for the project.
3. In the left panel, select **Build Path** and then select the **Project Libraries** tab.
4. Click **Add Library** and browse to and select the desired project library (.projlib) file.
5. Click **OK**. You are prompted to rebuild the project.
6. Save the resource. The project library appears at the root of the project tree as a Project Dependencies node.

To Import a Project Library

1. Do one of the following:
 - From the File menu select **Import**.
 - Right-click anywhere in BusinessEvents Studio Explorer and select **Import**.

- In the Import wizard Select dialog, select **TIBCO BusinessEvents > Project Library** and click **Next**. You see the Import Project Library dialog.



- In the Import Project Library dialog Select File field, browse to and select the project library (.projlib) file.
- Select a project from the project tree or type the name in the field above the project tree area. You can only import at the root of a project.
Ignore the New File Name field. It is not used by TIBCO BusinessEvents.
- Click **Finish**. The project library appears at the root of the project tree as a Project Dependencies node.

To Remove a Project Library

To remove a project library, follow the instructions in [To Include a Project Library \(Add it to the Build Path\)](#), page 12, but instead of clicking Add Library, select the project library you want to remove and click Remove Library.

Working with Global Variables

Global variables provide an easy way to set defaults for use throughout your project. When the project is deployed all occurrences of the global variable name are replaced with the provided global variable value, or a deploy-time override.

For example, you could assign the value 7474 to the global variable `RvDaemon`. You can then use the variable in a Rendezvous Transport resource. At deploy time you can override the default value as needed.



- The datatype of the global variable must match the datatype accepted in the field where you use it. If the global variable is of a different type, runtime errors result.
- An exception to the above allows flexibility in numeric fields: global variables used in numeric fields can be of any type, as long as the substituted value of the field is numeric.
- A project folder called `defaultVars` is available but not exposed in BusinessEvents Studio Explorer, so that you can share the global variables using source control software. It is not used for other purposes.



Setting and Overriding Global Variables from a Project Library

You can export global variables to a project library and import the library into other projects (see [Working with Project Libraries on page 10](#)).

When multiple global variables have the same name, one overrides the rest and is used in the project. The name, datatype, and default value of the overriding global variable are used. The override order is as follows:

- When multiple project libraries have the same-named global variable, the global variable from the library closest to the top of the project library list overrides any in lower libraries.
- When a local project global variable has the same name as a project library global variable, the local project global variable overrides the project library global variable.

The Global Variables view presents a merged list. A blue upward pointing arrow in the row for a global variable () indicates that this global variable overrides another global variable. The source of the global variable is shown in the Project Source column to the right of the arrow.

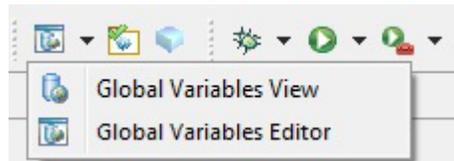
To see all global variables, including overridden ones, open the global variable editor and expand the Project Dependencies and project libraries. Global variables from project libraries are not editable.

Adding and Managing Global Variables

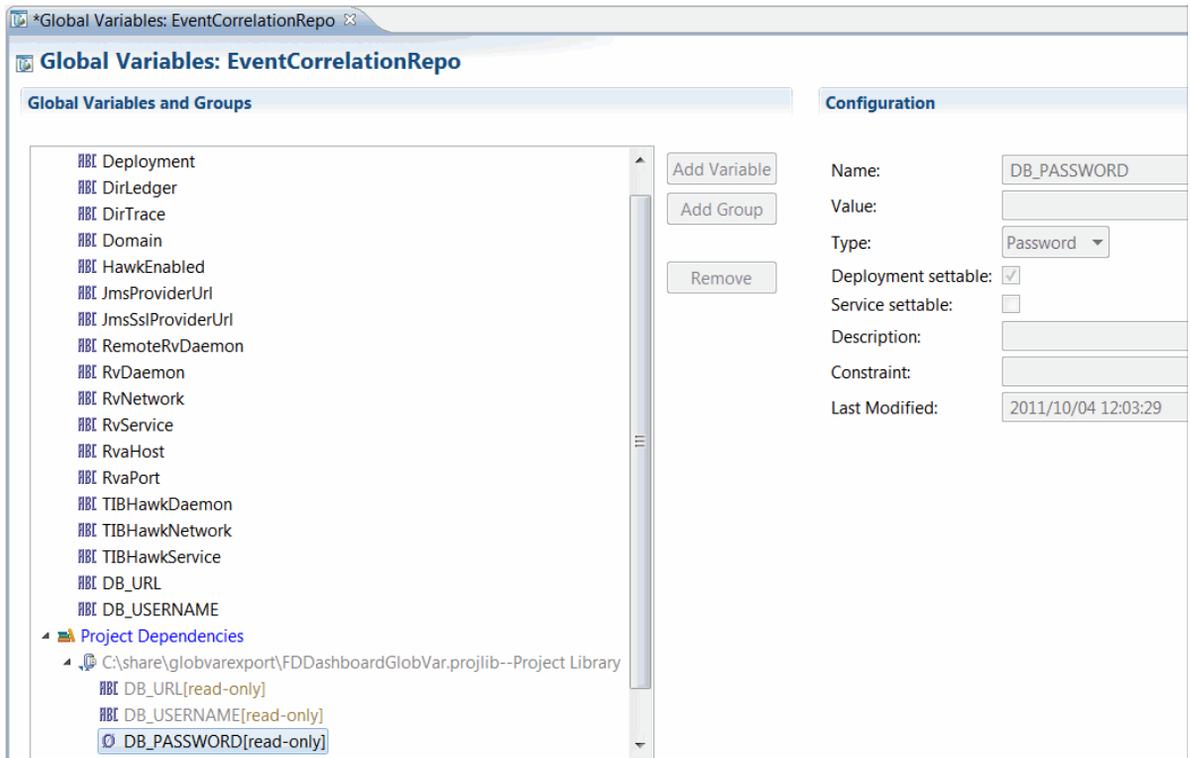
In a new TIBCO BusinessEvents Studio project no global variables are predefined. When you import a project from an earlier version of TIBCO BusinessEvents, however, you see predefined global variables (as well as any others defined in the project).

You may also see global variables in project libraries. All global variables are visible in the editor including overridden ones. To see the merged list of global variables that are used in the project, open the Global Variables view in one of the following ways:

- Click the **Global Variables View** button from the toolbar.



- From the top menus, select **Window > Show View > Other > TIBCO BusinessEvents > Global Variables**.



To Add and Manage Global Variables

1. Open the project in TIBCO BusinessEvents Studio.
2. Open the Global Variables Editor in one of the following ways:
 - From the toolbar, click the **Global Variables Editor** button.
 - From the top menus, select **Project > Edit Global Variables**.
 You see the Global Variables Editor listing the variables available, if any.
3. Do any of the following. (See [Table 3, Global Variable Reference](#), on page 16 for a guide to the fields):
 - To add a variable, click **Add Variable** and complete the fields.
 - To edit a variable, select the variable and update the fields.
 - To add a variable group, click **Add Group**.
 - To add a variable to a group, first select a group, then click **Add Variable**.
 - To remove a variable or a group, highlight it and click **Remove**.



Using Groups Groups are used for organizing variables. Variable groups are especially useful if multiple developers share a project using a version control system. When referencing a variable that is in a group, use the complete path, for example `%%mygroup/mysubgroup/myvariable%%`. (Because the complete path is used, the name of a variable in a group can be the same as the name of a variable in a different group.)

You must add at least one variable to a group, or the group will not be saved. If you delete all global variables in a global variable group, the group itself is also automatically deleted.

4. Save the resource. Groups and references to the `defaultVars.substvar` file appear in the `defaultVars` project folder.

Table 3 Global Variable Reference

Field	Description
Name	The variable name. Different groups can have the same-named variable, because the group name is included when you use the global variable.
Value	The variable value. Varies according to type.

Table 3 Global Variable Reference

Field	Description
Type	<p>One of String, Integer, Boolean, Long, Password.</p> <p>The type must match the type of the field where the global variable is used, or errors result.</p>
Deployment settable	<p>For deployment with TIBCO BusinessEvents Monitoring and Management You can only override global variables if both the Deployment settable and Service Settable checkboxes are checked.</p> <p>For deployment with TIBCO Administrator If checked, the variable is visible and settable when deploying using TIBCO Administrator. The values set at that time are saved in the project that TIBCO Administrator creates from the provided EAR file.</p> <p>If the checkbox is not checked, the variable is not visible in TIBCO Administrator. It is checked by default.</p>
Service Settable	<p>For deployment with TIBCO BusinessEvents Monitoring and Management You can only override global variables if both the Deployment settable and Service Settable checkboxes are checked.</p> <p>For deployment with TIBCO Administrator If both Deployment Settable and Service Settable are checked, the value of the global variable can be set differently for each deployable instance.</p> <p>Note Even if Service Settable is checked, the variable is included in the EAR only when the Include all service level global variables option is selected when building the enterprise archive file.</p> <p>In all deployment scenarios, values set at the service instance level are passed to the engine at runtime in the format <code>tibco.clientVar.VariableName=value</code>.</p> <p>Not checked by default.</p>
Description	A helpful description, as needed.
Constraint	<p>Optional. For String and Integer types, allows you to provide a range of allowed values. The constraint field for Strings is an enumeration, for example, <code>one</code>, <code>two</code>, <code>three</code>. The constraint field for Integers is for a range, for example, <code>1-100</code>.</p> <p>Note that constraints for Integers are currently not implemented in TIBCO Administrator.</p>
Last Modified	Non-editable field records date and time this variable was last modified.

Using Global Variables

To Use Global Variables in TIBCO BusinessEvents Studio Project Fields

To use a global variable as the value for a project setting, drag it from the list of variables into the text box for the setting, or enter it manually. Use the following syntax:

```
%%Variable_Group/Variable_Name%%
```

As shown above, you must include the global variable group hierarchy, if one exists.

For example, to use a global variable in a File Path field, you might enter the following:

```
%%filePathVars/certificateFilePath%%
```

To Use Global Variables in the Rule Editor

To use a global variable in the rule editor, use one of the `System.getGlobalVariableAs*` functions. For example:

```
System.getGlobalVariableAsString("myvars/Hostname", "Localhost")
```

Where `myvars/Hostname` is the name of the variable group and variable, and `Localhost` is an optional literal value to use if the variable is not found.

To Use Global Variables in Debugger

To use a global variable in Debugger, add it as a VM argument. Prefix the variable with `-v`, as shown:

```
-vVariable_Group/Variable_Name=
```

Overriding Global Variables at Deploy Time

You can override default values by setting global variable values in one of these ways, depending on how you will deploy:

- For starting at the command line:
 - In the design time CDD file. See [Setting Global Variables in the CDD File \(for Command Line Startup\) on page 402](#).
 - At the command line, using the `--propVar` option, or using the `-p` option to specify a property file where the override properties are defined. See [Starting a TIBCO BusinessEvents Engine at the Command Line in TIBCO BusinessEvents Administration](#).

- For deployment using TIBCO BusinessEvents Monitoring and Management component, set values in the CDD or MM Console. CDD values can be overridden in MM Console. See *Setting Global Variables in MM in TIBCO BusinessEvents Administration*.
- For deployment using TIBCO Administrator, set overrides in the TIBCO Administrator UI. You can override variables at the deployment level or at the service level. Values set at the service level are used for the specific engine you are deploying. Service settable global variables are only available if the Include All Service Level Global Variables checkbox in the Build Enterprise Archive dialog is checked (see [Building an Enterprise Archive \(EAR File\) on page 22](#)). Also see *Deploying a Project in a TIBCO Administrator Domain in TIBCO BusinessEvents Administration*.

Order of Precedence of Global Variable Overrides

Global variable values are selected at runtime using values set in the following ways, shown in the order of precedence, highest to lowest:

1. Command-line arguments at engine startup (highest priority).
2. Property files specified at command-line engine startup
3. MM Console or TIBCO Administrator
4. CDD file (ignored by TIBCO Administrator)
5. TIBCO BusinessEvents Studio Global Variable Editor (lowest priority)

Storing Trusted Certificates Outside of Your Project

Trusted certificates are used when you configure SSL, such as in an [HTTP Connection](#) or [JMS Connection](#).

Trusted certificates can be used to ensure that remote servers are who they claim to be and to ensure that TIBCO BusinessEvents can identify itself as a valid client when connecting to a server.

You can store the certificates within a project folder, or you can use a special global variable, `BE_GLOBAL_TRUSTED_CA_STORE`, to specify the location of an external directory that contains all the certificates known to TIBCO BusinessEvents.

When you store the certificates within a project folder, then when a certificate changes or expires, you must import any new certificates or certificate chains into the project, rebuild the EAR file, and re-deploy your project.

Using the global variable, however, avoids this problem. When you use the global variable to specify the external location of certificates, then when certificates change or expire, you just replace certificates or add new certificates and then restart the engine to load the changes.

You can set the global variable value and then use the variable in the usual ways, as described in this chapter. For example you could use the global variable as follows:

```
tibco.clientVar.BE_GLOBAL_TRUSTED_CA_STORE=file:///somePath/myGTCAFolder
```

To Store Trusted Certificates Outside of the Project

1. Create a directory where you want to store the trusted certificates. You must copy this directory to each machine where engines are deployed. Alternatively, the location can be a shared network area accessible by all process engines.
2. Create a global variable named `BE_GLOBAL_TRUSTED_CA_STORE`. See [Working with Global Variables on page 14](#) for more information.
3. Set the value of `BE_GLOBAL_TRUSTED_CA_STORE` to the location of the trusted certificates folder on your file system. The value must be a file URL, for example, `file:///c:/tibco/certs`.

The location can be the same for all deployed engines (that is, you copied it to the same location on each machine or it is a shared network drive).

Alternatively you can change the value of the global variable as needed when you deploy the project

4. Specify a value in the Trusted Certificates field in the SSL Configuration dialog. When the project runs, the value of `BE_GLOBAL_CA_STORE` is used, and not the value you specify in the Trusted Certificates field.
5. Save the resource.

Building an Enterprise Archive (EAR File)

You can build an enterprise archive file using a TIBCO BusinessEvents Studio dialog, and also using a command-line utility (see [Building an Enterprise Archive \(EAR File\) at the Command Line on page 26](#) for details).



For deployment using TIBCO Administrator, the configuration Name field value must match the project name The project does not deploy if they are different. The actual EAR file name, however, can differ from the configuration name.

Do not store the EAR file in a project folder Doing so includes the prior EAR file when you build the EAR file again, needlessly increasing the size.

Modifying Ignored Resources List Certain files (and folder names) are excluded from the EAR. To maintain the list of exclusions, in TIBCO BusinessEvents Studio, select Window > Preferences > TIBCO BusinessEvents > Code Generation > Ignored Resources. See [Setting Preferences on page 726](#) for more on this topic.

EAR File Encoding

Default Encoding

The default encoding of the EAR files generated by TIBCO BusinessEvents Studio is ISO8859-1. This is also the default encoding of TIBCO Administrator. To upload an EAR file to TIBCO Administrator, the EAR file encoding must match the TIBCO Administrator encoding.

To Change the Default EAR File Encoding

To change the default EAR file encoding, define a global variable named `MessageEncoding` and set its value to the desired encoding. For example, UTF-8.

Options for Deploying the EAR

For deployment options see Chapter 5, *Deploying and Managing Engines with MM* and Chapter 8, *Deploying With TIBCO Administrator in TIBCO BusinessEvents Administration*. For testing purposes, you may find it useful to refer to the section *Starting a TIBCO BusinessEvents Engine at the Command Line in TIBCO BusinessEvents Administration*.

Building an EAR File in TIBCO BusinessEvents Studio

When testing project before development, you may want to build the EAR as explained here, and run the project at the command line.

By default, EAR files are built in memory. The compiler does not use the file system during code generation. Instead, the Studio JVM is used to load all the Java classes and resources into memory until the build process is completed. However, you can choose to use the file-system based compiler to build EAR files.

1. In TIBCO BusinessEvents Studio, select the menu **Window > Preferences** to open the Preferences dialog.
2. Select **TIBCO BusinessEvents > Code Generation** on the left panel and then set the **Compilation Mode** to **File System**.

When using the Studio Tools utility to build an EAR file, set the option `-lc` to use the file-system based legacy compiler.



When building an EAR file in memory for a large project, the JVM may run out of PermGenSpace and/or heap space. In such cases, edit the `BE-HOME/studio/eclipse/studio.ini` and `BE-HOME/studio/bin/studio-tools.tra` file to set appropriate values for the JVM settings. By default the heap size is set to `-XX:MaxPermSize=256m`.

1. In BusinessEvents Studio Explorer, highlight the project name, then from the top menus select **Project > Build Enterprise Archive**.
If you see a message asking you to save all project resources, click Yes. It means an unsaved resource editor is open.
2. At the Build Enterprise Archive dialog, complete values according to guidelines provided in [Table 4, Build Enterprise Archive Reference](#).
3. Click **Apply** to save the configuration details.
(To revert to the version already saved, click Revert.)
4. Click **OK** to build the archive.

Table 4 Build Enterprise Archive Reference

Field	Description
Name	<p>Name of this EAR configuration. (Not the EAR filename.)</p> <p>Default value is the project name.</p> <p>Note For deployment using TIBCO Administrator, the configuration Name field value must match the project name. The project does not deploy if they are different. The actual EAR file name, however, can differ from the configuration name.</p>
Author	<p>Person responsible for the EAR file.</p> <p>Default value is the currently logged-on user name.</p>
Description	Optional description.
Archive Version	Increments on each build of the EAR. You can also manually enter a version identifier.
Generate Debug Info	<p>Check this checkbox if you want to use the debugger. (See Chapter 40, Testing and Debugging Projects, on page 679.)</p> <p>Default setting is checked.</p>
Include all service level global variables	Check to include service level global variables.
File Location	Browse to the directory in which you want to store the EAR and enter an EAR filename.
Delete Temporary Files	<p>Before TIBCO BusinessEvents packages the EAR file, it generates the Java code in a temporary directory. After the files are packaged in the EAR file, then the temporary files and directory are deleted.</p> <p>You can choose to keep the generated Java files, for example to troubleshoot some problem with an EAR file. To do so, uncheck the Delete Temporary Files checkbox, and specify where to store the Java files in the Compilation Directory field.</p> <p>Default setting is checked, meaning that temporary files are not saved.</p>
Compilation Directory	If you uncheck the Delete Temporary Files checkbox, specify the directory where you want to save the Java files generated during the process of building the EAR file.

TIBCO BusinessEvents Studio Tools is a command-line utility with various operations (tools) you can use to automate common procedures. This chapter documents the tools that can be useful while developing applications in TIBCO BusinessEvents Studio. Other tools in this suite are documented in *TIBCO BusinessEvents Decision Manager User's Guide* and *TIBCO BusinessEvents Installation*



All arguments of a command-line utility must use only ASCII characters.

Topics

- [Building an Enterprise Archive \(EAR File\) at the Command Line, page 26](#)
- [Importing an Older Studio Project at the Command Line, page 28](#)
- [Working With Project Libraries at the Command Line, page 31](#)
- [Migrating Core Coherence Functions at the Command Line, page 34](#)
- [Generating All Project Class Files at the Command Line, page 35](#)
- [Generating Encrypted Passwords, page 37](#)
- [Running the Tester on the Command Line, page 38](#)

Building an Enterprise Archive (EAR File) at the Command Line

For full details on deployment, see *TIBCO BusinessEvents Administration*. The `buildEar` operation within the `studio-tools` utility is useful for automation purposes, for example, in testing environments.

By default, the EAR files are built in memory. The compiler does not use the file system during code generation. Instead, the Studio JVM is used to load all the Java classes and resources into memory until the build process is completed. You can choose to use the file-system based compiler to build EAR files by setting the appropriate options.

To Build an EAR File at the Command Line

1. Navigate to `BE_HOME/studio/bin/` and open a command prompt.
2. Execute a command with the following format (all on one line) at a command prompt:

```
studio-tools -core buildEar [-h] [-x] [-lc] [-o outputEarFile>] -p
studioProjectDir [-pl projectLibrariesPath] [-cp extendedClasspath]
```

For example:

```
studio-tools -core buildEar -o c:\FD.ear -p
D:\Workspace\FraudDetection
```

[Table 5, TIBCO BusinessEvents Studio Tools Options for Building an EAR File](#), provides detailed information about the options.

Table 5 TIBCO BusinessEvents Studio Tools Options for Building an EAR File

Option	Description
<code>-core buildEar</code>	Within the <code>core</code> category of operations, specifies the <code>buildEar</code> operation for building EAR files.
<code>-h</code>	Optional. Displays help.
<code>-x</code>	Optional. Overwrites the specified output file if it exists.
<code>-lc</code>	Optional. Specifies that the file-based legacy compiler must be used to build the EAR file. By default, the EAR files are built in memory.
<code>-o</code>	Optional. Specifies the filename for the output EAR file. If not specified the EAR file is the same as the final (leaf) directory name in the <code>projectDir</code> path.

Table 5 TIBCO BusinessEvents Studio Tools Options for Building an EAR File (Cont'd)

Option	Description
-p	Absolute path to the TIBCO BusinessEvents Studio project directory. The EAR file is built using this project.
-pl	Optional. Specifies list of project library paths to be used separated by a path separator
-cp	Optional. Specifies the extended classpath to be used.



When building an EAR file in memory for a large project, the JVM may run out of PermGenSpace and/or heap space. In such cases, edit the *BE-HOME/studio/eclipse/studio.ini* and *BE-HOME/studio/bin/studio-tools.tra* file to set appropriate values for the JVM settings. By default the heap size is set to `-XX:MaxPermSize=256m`.

Importing an Older Studio Project at the Command Line

This tool is equivalent to the following menu option in TIBCO BusinessEvents Studio: **File > Import > TIBCO BusinessEvents > Existing TIBCO BusinessEvents Studio Project**. It is used to migrate older(4.x and 5.0.x) projects to 5.1.

When you run this command-line utility, various migration actions are done, as explained in *Migration from 4.x to 5.1 in TIBCO BusinessEvents Installation*.

For migration of version 3.x projects, see Chapter 4, *Migrating Projects from Earlier Versions in TIBCO BusinessEvents Installation*.

Before you can use a 4.x project imported at the command-line in TIBCO BusinessEvents Studio you must do another procedure, explained in [To Open a Project Imported at the Command Line in TIBCO BusinessEvents Studio on page 29](#).

To Import an Existing 4.x Project at the Command Line

1. Navigate to `BE_HOME/studio/bin/` and open a command prompt.
2. Execute a command with the following format (all on one line) at a command prompt:

```
studio-tools -core importExistingProject [-h] -p studioProjDir [-o targetProjDir] [-c CDDprojectPath] [-u PUNameFromCDD]
```

For example:

```
studio-tools -core importExistingProject -p C:\FT\SomeProj -o c:\MyWorkspace\SomeProj -c COM.cdd -u Invproc
```

If HTTP channel properties are migrated (from a specified CDD and processing unit to all HTTP channel resources' Advanced tab) you see a message like this:

```
Migrating HTTP properties of Processing Unit "PUName" from CDD "CDDprojectPath" to HTTPChannel(s) present in the project
```

When the import has completed successfully, you see a message in the command window like the following:

```
The existing 4.0 TIBCO BusinessEvents project has been successfully imported to c:\MyWorkspace\SomeProj.
```

[Table 6, TIBCO BusinessEvents Studio Tools Options for Importing an Existing Project](#), provides detailed information about the options.

Table 6 TIBCO BusinessEvents Studio Tools Options for Importing an Existing Project

Option	Description
<code>-core importExistingProject</code>	Specifies the <code>importExistingProject</code> operation for importing a TIBCO BusinessEvents Studio project into the workspace.
<code>-h</code>	Optional. Displays help.
<code>-p</code>	Source project: absolute path to the project directory of the TIBCO BusinessEvents Studio project to be imported.
<code>-o</code>	<p>Optional. Absolute path to the target project directory, where the project is imported to.</p> <p>If you specify the source project directory name as the last element in the path, it is used as the target project directory. If you specify a different directory as the last element in the path, the directory is created if it does not exist, and the source project directory is imported within the specified target directory.</p> <p>If you do not specify a target project directory, the original project contents are updated. If the project to be imported is a TIBCO BusinessEvents version 4 project, it is no longer compatible with version 4 after the import.</p> <p>If the target location points to an existing project, the import does not proceed and this message displays:</p> <p>The specified target location already exists and cannot be used.</p>
<code>-c</code>	Optional. The CDD to use for migration actions. Project path of the CDD (path relative to the root directory of the source project).
<code>-u</code>	<p>Optional but if specified <code>-c</code> must also be specified. Specifies the name of the PU (within the specified CDD) that contains settings to be migrated.</p> <p>HTTP channel settings from this PU are migrated to all HTTP channel resources in the project.</p>

To Open a Project Imported at the Command Line in TIBCO BusinessEvents Studio

To open a project imported at the command line, you must add it as a new project.

1. Start TIBCO BusinessEvents Studio. In Windows, click **Start > All Programs > TIBCO > ENV_NAME > TIBCO BusinessEvents 5.1 > TIBCO BusinessEvents Studio**.
2. From the File menu select **New > Project**. You see the New Project — Select a Wizard dialog.
3. Select **TIBCO BusinessEvents > Studio Project** and click **Next**.
4. In the Project Name field, enter the directory name where the imported project is located. (This is used as the project name.)
5. (If you imported the project to a directory in your default workspace, skip this step.) If the project directory is located outside the default workspace, uncheck the Use default location checkbox and browse to the *parent* directory of the project imported at the command line.
6. Click **Finish**. The project folders appear in the Studio Explorer view.

Working With Project Libraries at the Command Line

See [Working with Project Libraries on page 10](#) for an introduction to project libraries and procedures for working with them in TIBCO BusinessEvents Studio.

This section explains how to use a command-line utility to do the following:

- Create a project library (and optionally overwrite any existing library at the same location with the same name). This action creates the `.projlib` file. The corresponding action in TIBCO BusinessEvents Studio is exporting a project library.
- Add an existing project library to a TIBCO BusinessEvents Studio project. The corresponding action in TIBCO BusinessEvents Studio is importing a project library (or adding it to the build path property page).
- Remove an existing project library from a TIBCO BusinessEvents Studio project. The corresponding TIBCO BusinessEvents Studio action is removing the library from the Build Path property page.

To Work with Project Libraries at the Command Line

1. Navigate to `BE_HOME/studio/bin/` and open a command prompt.
2. Execute a command with the following format (all on one line) at a command prompt. Arguments for creating a project library are shown separately for clarity:

To create a project library:

```
studio-tools -core buildLibrary [-h] -p studioProjDir [-a] [-r] [-x] -n
projectLib [-f resources]
```

To add or remove an existing project directory in a TIBCO BusinessEvents Studio project:

```
studio-tools -core buildLibrary [-h] -p studioProjDir [-a] [-r] [-x] -n
projectLib [-f resources]
```

Examples for Creating a Project Library

Example 1 The following command creates a project library using the specified project's contents:

```
studio-tools -core buildLibrary -p C:\workspace\MyProj -n
C:\test\myproj.projlib
```

Example 2 The following command creates a project library at the location specified, and overwrites any project library already at that location. From the specified project, the project library uses only the resources specified in the `-f` argument: the `Concepts` folder, the `MyRule.rule` and all dependent resources referred to by these resources.

```
studio-tools -core buildLibrary -p C:\workspace\MyProj -x -n
C:\test\myproj.projlib -f Concepts,Rules\MyRule.rule
```

Examples for Adding and Removing a Project Library

When a project library exists, built either using TIBCO BusinessEvents Studio or using the command-line options shown in [Examples for Creating a Project Library, page 31](#), you can use the `buildLibrary` operation to add a project library to or remove a project library from a TIBCO BusinessEvents Studio project.

Example 3 The following command adds the `myproj.projlib` project library to the TIBCO BusinessEvents Studio project at `C:\workspace\MyProj`.

```
studio-tools -core buildLibrary -p C:\workspace\MyProj -a -n
C:\test\myproj.projlib
```

Example 4 The following example removes the `myproj.projlib` project library from the TIBCO BusinessEvents Studio project at `C:\workspace\MyProj`.

```
studio-tools -core buildLibrary -p C:\workspace\MyProj -r -n
C:\test\myproj.projlib
```

[Table 7](#) provides detailed information about the options.

Table 7 TIBCO BusinessEvents Studio Tools Options for Working with Project Libraries

Option	Description
<code>-core buildLibrary</code>	Specifies the <code>buildLibrary</code> tool for adding, removing, and creating (overwriting as needed) project libraries.
<code>-h</code>	Optional. Displays help.
<code>-p</code>	File path to the TIBCO BusinessEvents Studio project. Resources from this project are used to create the project library. This filepath is also used to identify the project to which you want to add a project library.
<code>[-r -a]</code>	Use one of these arguments as needed. <code>-r</code> : removes the specified project library from the project specified by <code>-p</code> . <code>-a</code> : adds the specified project library to the specified project.

Table 7 TIBCO BusinessEvents Studio Tools Options for Working with Project Libraries (Cont'd)

Option	Description
-x	Optional. Overwrites any existing project library at the location specified by -n.
-n	The absolute path of the project library to be created or added or removed
-f	Optional. A comma-separated list of resources to include in the project library, using a path relative to the <code>projectDir</code> . If not specified all resources of the specified project are used.

Migrating Core Coherence Functions at the Command Line

The TIBCO BusinessEvents DataGrid component is the default cache provider for TIBCO BusinessEvents. Use of Oracle Coherence is also supported for those who want to use their own copy of this software.

A set of core Coherence functions was renamed in TIBCO BusinessEvents 5.0 (and additional internal changes were made) so that these functions can be used with either the Coherence or the TIBCO BusinessEvents DataGrid cache provider.

Migrating existing projects to use the new function names is one of the actions done automatically when you import a project from a previous release, as explained in Chapter 4, *Migrating Projects from Earlier Versions in TIBCO BusinessEvents Installation*.

Separate use of the `migrateCoherenceCalls` operation may be required in some cases. For example, you create a project in the current version with Coherence as the cache provider, and enable and use the Coherence-specific catalog functions. Then you switch to using TIBCO BusinessEvents DataGrid as the cache provider. You would then use the `migrateCoherenceCalls` operation to migrate those functions.

To Migrate Core Coherence Functions at the Command Line

1. Navigate to `BE_HOME/studio/bin/` and open a command prompt.
2. Execute a command with the following format (all on one line) at a command prompt:

```
studio-tools -core migrateCoherenceCalls -projectDir studioProjectDir
```

When the migration has completed successfully, you see a message in the command window:

```
All Coherence function calls have been migrated successfully.
```

Generating All Project Class Files at the Command Line

You can generate all class files in a project at the command line. Although this is a core component, the class files are generally used within the context of TIBCO BusinessEvents Decision Manager, where decision table class files can be separately deployed.

To Generate Class Files at the Command Line

1. Navigate to `BE_HOME/studio/bin/` and open a command prompt.
2. Execute a command with the following format (all on one line) at a command prompt:

```
studio-tools -core generateClass [-h] -p studioProjectDir [-n
studioProjectName] -o outputPath [-x {true | false}] [-lc] [-pl
projectLibraryPath] [-cp extendedClasspath]
```

For example:

```
studio-tools -core generateClass -p D:\Workspace\FraudDetection -o
c:\temp -x true -cp c:\tibco\be\5.1\lib\myjar.jar
```

You see a success message if the files were generated successfully.

[Table 8, TIBCO BusinessEvents Studio Tools for Generating Class Files](#), provides detailed information about the options.

Table 8 TIBCO BusinessEvents Studio Tools for Generating Class Files

Option	Description
<code>-core generateClass</code>	Within the <code>core</code> category of operations, specifies the <code>generateClass</code> operation used to generate a project's class files.
<code>-h</code>	Optional. Displays help.
<code>-p</code>	Absolute path to the TIBCO BusinessEvents Studio project directory.
<code>-n</code>	Optional. Specifies the name of the TIBCO BusinessEvents Studio project whose class files are to be generated. If not specified, the final (leaf) directory name in the path specified for the <code>-p</code> option is used as the project name.

Table 8 TIBCO BusinessEvents Studio Tools for Generating Class Files (Cont'd)

Option	Description
-o	<p>Specifies the output directory for generated classes.</p> <p>If you do not specify a directory, files are placed in a user temporary directory. For example, on Windows files might go in a directory like the following:</p> <pre>C:\Documents and Settings\User\Local Settings\Temp\BE_1322046141896</pre>
-x	Optional. If <code>true</code> , overwrites any existing class file with the same name.
-lc	Optional. Specifies that the file-based legacy compiler must be used to build the EAR file. By default, the EAR files are built in memory.
-pl	Optional. Specifies list of project library paths to be used separated by a path separator.
-cp	<p>Optional. Extended classpath. Use as needed. Provide separate JAR file paths for each JAR file required for project compilation. For example, additional classpath information is needed if the decision table uses custom functions or third-party JAR files.</p> <p>Separate entries by the appropriate path separator. For example if the separator is semicolon (;) you might add the following:</p> <pre>C:\customjars\custom.jar;C:\customjars\custom2.jar</pre>

Generating Encrypted Passwords

TIBCO BusinessEvents Studio Tools utility provides a command to generate encrypted passwords.

Syntax `studio-tools -crypto encrypt [-h] [-i inputText]`

Definition The encrypted text is written to STDOUT. The `exit` status of the command indicates whether the generation was a success or a failure.

An `exit` status of `0` indicates success, and `-1` indicates an error.

[TIBCO BusinessEvents Studio Tools Options for Generating Encrypted Passwords on page 37](#) provide more detailed information about the option.

Table 9 TIBCO BusinessEvents Studio Tools Options for Generating Encrypted Passwords

Option	Description
-h	Optional. Displays help.
-i	Input text which needs to be encrypted

Running the Tester on the Command Line

Running the Tester on the command line allows you to create test data templates, populate the test data files, and run tests.

To run the Tester on the command line, you will follow these steps:

1. Run the command

```
-tester generate
```

This generates the test data template files, which are used for concepts, scorecard, and events for a given project directory.

The test data template files are either in XLS or CSV format, as desired by the user, and are used to enter the test data. They will have a project name, element relative path, external ID, column names, and so on.



Blank columns should be filled with NA.

To generate a file in a specific format, see [Command Line Options, page 39](#).

2. Enter the data manually in the generated template files.
3. Run the command

```
-tester import
```

This will import the data entered into test data files

```
* .concepttestdata
```

and

```
* .eventtestdata.
```

These files are similar to the files saved by the Studio Tester when selecting Create Test Data.

4. Run the command

```
-tester assert
```

This will start an instance of the TIBCO BusinessEvents engine and assert the test data generated in step 2 to the instance.

After the assertion, the result data files will be saved in an output directory that has to be provided in the user command line options, and in one of these two formats:

- **Default Format:** It is either the XML format (default) or HTML format (defined through the command line option).
- **Custom Format:** Output is based on Types or Operation (such as Created, Modified, and Deleted).

The default output is generated based on Operation.

Command Line Options

The `-tester` operation provides for the following commands:

- `generate`
- `import`
- `assert`

generate

The command `generate` reads concepts, scorecards and events from a project and creates test data template files in comma separated format. The output files can be edited by the user to add test data and then can be used as an input for the `import` command to import the test data.

Syntax `studio-tools -tester generate [-h] [-x] -p <projectDir> [-t <xls | csv >] [-s <separator>] [-o <outputDirectory>]`

Options The following options are used with the command `generate`:

Table 10 Options for the Command `generate`

Option	Description
<code>-h</code>	(optional). Prints usage help
<code>-o</code>	(optional). Specifies the output directory. If not specified, the output directory will be <code>ProjectDirectory/TestDataTemplates</code> .
<code>-p</code>	(required). Specifies the project path.
<code>-t</code>	optional). Specifies the output file type. Supported values are [<code>xls</code> <code>csv</code>]. If this option is not specified, the output will be in CSV format with a comma (,) as a column separator.
<code>-s</code>	(optional. Used only when the output file type is CSV to specify a column separator. The supported values are [<code>,</code> <code>;</code>]. If not specified, the default column separator is comma (,).
<code>-x</code>	(optional). Overwrites the output directory if it already exists.

import

The command `import` reads the CSV data files and converts them to the appropriate test data format (such as `.concepttestdata`, `.eventtestdata`).

Syntax `studio-tools -tester import [-h] [-x] -i <inputFile> -p <projectDirectory> [-o <outputDirectory>] [-t < xls | csv >] [-s <separator>]`

Options The following options are used with the command `import`:

Table 11 Options for the Command `import`

Option	Description
-i	(required). Sets the input path, which is a path to a directory containing test data template files, a single test data template file, or a path to <code>Listfile</code> containing paths of files that must be included for test data generation.
-h	(optional). Prints usage help
-o	(optional). Specifies the output path.
-p	(required). Specifies the project path.
-t	optional). Specifies the output file type. Supported values are [<code>xls</code> <code>csv</code>]. If this option is not specified, the output will be in CSV format with a comma (,) as a column separator.
-s	(optional. Used only when the output file type is CSV to specify a column separator. The supported values are [<code>,</code> <code>;</code>]. If not specified, the default column separator is comma (,).
-x	(optional). Removes the <code>TestData</code> root from under the output directory, if it already exists, and creates a new one.

assert

The command `assert` is used to assert the test data on a running TIBCO BusinessEvents engine instance.

Syntax `studio-tools -tester assert [-h] -p <projectDir> -c <cddFile> -r <earFile> -w <workingDir> [-u <processingUnit>] [-f <outputFormatType>] [-o <outputDirectory>] [-n <filename>] [-i <input path>] [-x]`

Options The following options are used with the command `assert`:

Table 12 Options for the Command `assert`

Option	Description
<code>-c</code>	(required). Specifies the CDD file path.
<code>-f</code>	(optional). Specifies the output format type. Acceptable values are <code>ops</code> and <code>type</code> . If this option is not specified, the default output is generated.
<code>-h</code>	(optional). Prints usage help
<code>-i</code>	(optional). Specifies the input directory path. If not specified, it will attempt to use the <code>TestData</code> directory under the project directory.
<code>-n</code>	(optional). Name of the <code>resultFile</code> . If not specified, it creates a file with the name <code>Run-#.resultdata</code> .
<code>-o</code>	(optional). Specifies the output directory. If not specified, the output directory will be <code>ProjectDirectory/results</code> .
<code>-p</code>	(required). Specifies the project path.
<code>-r</code>	(required). Specifies the EAR file path.
<code>-u</code>	(optional). Specifies the processing unit name. If not specified, it will use the default processing unit.
<code>-w</code>	(required). Specifies the working directory path.
<code>-x</code>	(optional). Overwrites the output file if it already exists.

Element Refactoring Operations

When you make changes to a project element, all references to that element must be updated accordingly. When such changes affect only the structure and not the behavior of the project, this operation is known as *project refactoring*. This chapter explains how to use the refactoring features, and some related features to do with copy-paste operations.

Topics

- [Renaming, Moving, Deleting, and Copy-Pasting Elements, page 44](#)
- [Migrating Core Coherence Functions, page 50](#)
- [Automatic Refactoring Actions and Limitations, page 53](#)

Renaming, Moving, Deleting, and Copy-Pasting Elements

Changes that affect the structure of a project, but not its behavior, are known as *project refactoring* changes. Refactoring ensures that the project structure remains self-consistent.

Copy-paste operations are not strictly speaking refactoring operations. However some refactoring is also done to support these operations, so they are included here.



You can copy items from one project to other projects in the workspace. Ensure that the items are suitable for their destination projects.

Moving, renaming, deleting, or copy-pasting project elements are changes that often affect other parts of a TIBCO BusinessEvents Studio project. Names of elements, element properties, and element locations, are referenced in various parts of a project such as rules, rule functions, and concept relationship properties.

When you make changes to a project element, all references to that element must be updated accordingly. The refactoring wizard has a preview page that enables you to review all these changes (see [Working with the Preview Page on page 48](#)).

This section explains the refactoring (and related) procedures. See [Automatic Refactoring Actions and Limitations on page 53](#) to understand what TIBCO BusinessEvents does for each type of refactoring operation.

Updating All References is Strongly Recommended

To ensure the integrity of the project, it is strongly recommended that you make all changes to all locations where a renamed or moved items is referenced. Only disable such changes if you are certain there are no references to the element, or there are unusual circumstances that justify such action.



Eclipse Tips

To Undo Changes You can undo additions or deletions in a project. Click **Edit > Undo** or press **Ctrl+Z**.

To Revert to an Earlier Version or Restore a Deleted Resource Eclipse also allows you to compare your work with local history and to revert to an earlier saved version (not just the last saved version). Right-click a resource in BusinessEvents Studio Explorer and select **Compare With > Local History** or **Replace With > Local History**. Also, if you right-click a folder and select **Restore From Local History** you can restore items deleted from that folder.

Project Level Actions

You can rename a project (use File > Rename) and you can copy and paste a project. However, you cannot move a project.

Renaming, Moving, and Deleting Elements

Renaming, moving, and deleting are refactoring actions that can have an effect on other parts of the project where element names and locations are referenced.

To Rename a Project Element

1. Rename the element using one of the following methods:
 - Right-click a project element in BusinessEvents Studio Explorer and select **Refactor > Rename**.
 - Highlight the element, then select **File > Rename**.
 - Highlight the element, then press **F2**.

The first page of the Rename Element wizard appears.

2. In the New Name field, type the new name.
3. If you do not want to rename the element in places where it is referenced, uncheck the Update References checkbox (but see [Updating All References is Strongly Recommended on page 44](#)).
4. Click **Preview**. One of the following occurs:
 - The preview page displays so you can examine the effect of this change. See [Working with the Preview Page on page 48](#) for details.
 - The problems page displays if the rename cannot be done, for example, because a new element name is used by an existing element. Click **Back** to fix the problem, or **Cancel** to cancel the rename.
5. In the preview page, clear checkboxes if you do not want the change to be made in some referenced locations. This is not generally recommended (see [Updating All References is Strongly Recommended on page 44](#)).
6. To complete the change, click **OK**.

To Move an Element to a Different Project Folder

1. Move the element using one of the following methods:
 - Right-click a project element name in BusinessEvents Studio Explorer and select **Refactor > Move**.
 - Highlight the element, then select **File > Move**.
 - Drag the element to the target folder (go to [step 3](#)).
2. If you opened the Move Element wizard using menus, navigate the project tree in the first page of the wizard to select a destination folder, then click **Preview**.
3. One of the following occurs:
 - The preview page displays so you can examine the effect of this change. See [Working with the Preview Page on page 48](#) for details.
 - The problem page displays if the move cannot be made. Click **Back** to fix the problem, or **Cancel** to cancel the move.
4. In the preview page, clear checkboxes if you do not want the change to be made in some referenced locations. This is not generally recommended (see [Updating All References is Strongly Recommended on page 44](#)).
5. To complete the change, click **OK**.

To Delete an Element or Folder

When you delete a folder, all elements within that folder are also deleted.

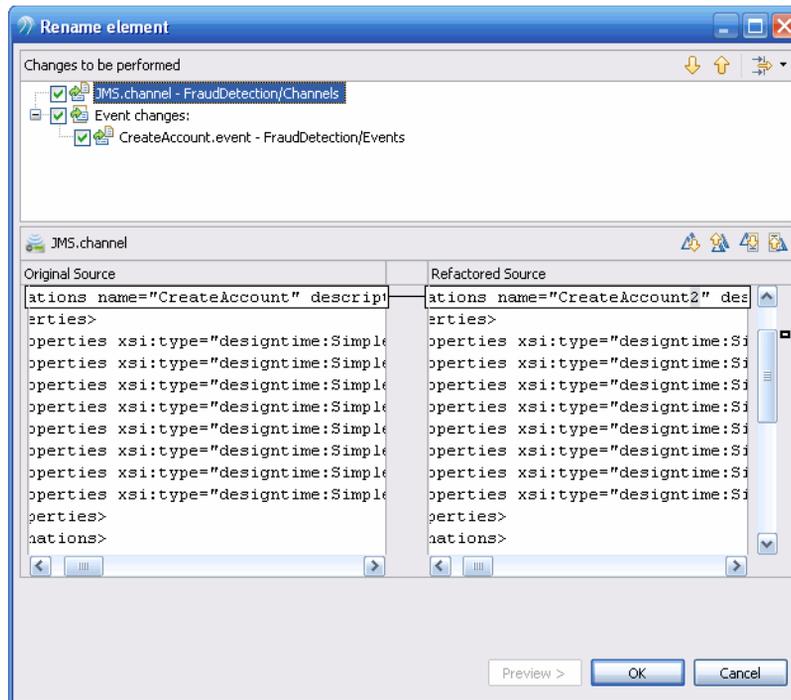
1. Do one of the following:
 - Right-click a project element or folder name in BusinessEvents Studio Explorer and select **Delete**.
 - Highlight the element, then select **Edit > Delete**.
 - Highlight the element, then press the **Delete** key.
2. At the Delete Resources page, click **OK** to delete without previewing, or click **Preview** to preview the effect of the deletion.
3. If you click Preview, one of the following occurs:
 - The preview page displays so you can examine the effect of this change. (For deletions, there is generally no information.) Click **OK** to delete.
 - The problems page displays if there is a problem with the deletion. Click **Back** to fix the problem, or **Continue** to force the deletion (or **Cancel** to cancel the deletion).

To Delete a Project

1. Do one of the following:
 - Right-click the project name in BusinessEvents Studio Explorer and select **Delete**.
 - Highlight the project name, then select **Edit > Delete**.
 - Highlight the project name, then press the **Delete** key.
2. If you want to remove the project contents on disk, check the Delete Project Contents on Disk check box. If you do not check this checkbox, then the project is removed from BusinessEvents Studio Explorer, but the project contents remain on disk.
3. Click **OK** to delete without previewing, or click **Preview** to preview the effect of the deletion.
4. If you click Preview, one of the following occurs:
 - The preview page displays so you can examine the effect of this change. (For deletions, there is generally no information.) Click **OK** to delete.
 - The problems page displays if there is a problem with the deletion. Click **Back** to fix the problem, or **Continue** to force the deletion (or **Cancel** to cancel the deletion).

Working with the Preview Page

During refactoring operations, you can click Preview to examine the effects of the operation on the project. (If a pre-check finds issues that may prevent the operation from completing successfully, a problem page appears instead.) This section explains how to use the preview page.



Checks in the upper panel indicate elements that will change as a result of a refactoring operation. Expand to take a closer look at individual folders and elements. All elements that appear are affected by the change. When you highlight an element, details of the change to be made are shown in the lower panel.

The Original Source panel on the left and the Refactored Source on the right use the persisted format of the element. Change bars indicate changed areas.

Do any of the following as needed to examine the changes and select a subset of the changes to be done on clicking OK:

- In the upper panel, use the arrows to navigate up and down a long list of changes.
- Click the Filter Changes button and select Hide Derived Resources. For example, a diagram is a derived resource. Diagrams are not persisted. You can easily recreate them. So you may not be interested in seeing those changes.

- If you want to apply the change to only some of the project elements, uncheck the checkboxes next to the elements as desired. For example, you may wish to replace or delete an element after you have completed the refactoring operation, so you do not need to apply the change to that element.



It is recommended that you accept all changes to be performed Your project can become corrupted if you do not make the changes throughout the project. Only deselect changes if you have a specific, valid reason to do so.

Copy-Pasting an Element

Copy-pasting is not a true refactoring operation, because it does change the behavior of the TIBCO BusinessEvents Studio project. You would generally make manual changes to your project to use the newly created project element. However, some limited refactoring is done for your convenience. For example, the definition of a rule or rule function begins with its fully qualified name, such as the following:

```
rule SomeRules.Application_Rule
```

If you copy and paste the above rule into the folder `OtherRules`, the definition of the rule automatically changes to:

```
rule OtherRules.Application_Rule
```

To Copy-Paste an Element

1. In BusinessEvents Studio Explorer, highlight the element (or folder) and do one of the following:
 - Right-click the element and select **Copy**.
 - Press **Ctrl+C**.
 - Select **Edit > Copy**.
2. Highlight the project folder where you want to paste the element, and do one of the following:
 - Right-click the folder and select **Paste**.
 - Press **Ctrl+V**.
 - Select **Edit > Paste**.

If you are pasting an element to the same folder that you copied it from, a dialog enables you to provide a different name. The default value is `CopyOfoldname`.

Migrating Core Coherence Functions

The TIBCO BusinessEvents DataGrid component is used as the default cache provider. Use of Oracle Coherence is also supported for those who want to use their own copy of this software.

A set of core functions has been renamed (and additional internal changes have been made) so that these functions can be used with either the Coherence or the TIBCO BusinessEvents DataGrid cache provider.

The renamed functions are in the Standard catalog `Cluster.DataGrid` category. You must use these functions going forward, and not the prior versions. See [Table 13, Mapping of Coherence Functions to DataGrid Functions, on page 51](#) for details.

For all projects created prior to version 5.0.0, use the 4.x project import wizard, which migrates the core Coherence function calls automatically.

For occasional use cases where the function calls are migrated manually, you can use the manual refactoring operation. You can also use a command-line option as explained in [Migrating Core Coherence Functions at the Command Line on page 34](#).

To Migrate Coherence Function Calls in TIBCO BusinessEvents Studio

See [Table 13, Mapping of Coherence Functions to DataGrid Functions, on page 51](#) for details of the functions that are mapped using this feature.

1. Open the project in TIBCO BusinessEvents Studio (3.x projects must be migrated first). Migrate the function calls using one of the following methods:
 - Right-click the project name in Studio Explorer and select **Refactor > Migrate Coherence Function Calls**.

The first page of the Migrate Coherence Function Calls wizard appears. It displays an informational list of the functions that the utility renames.

2. Click **OK** to make all necessary changes without previewing, or click **Preview** to preview the effects of the change.
3. If you click Preview, one of the following occurs:
 - The preview page displays so you can see the list of affected elements (which include rules, rule functions, expiry actions and state modeler transition, exit, and entry rules). You can scroll the list and view the original and refactored sources for each affected element. You can uncheck the checkboxes of elements you wish to remain unchanged (and later make

changes manually as needed). Click **OK** to make the changes to all or selected elements.

- The problems page displays if there is a problem with the change. Click **Back** to fix the problem, or **Continue** to force the change (or **Cancel** to cancel the change).
- The preview page displays this message: The refactoring does not change any source code. Click **OK** and complete the migration anyway.



Complete the refactoring operation even if no project code is affected This refactoring operation makes additional changes to your projects so that your projects will work correctly with the new function names.

Mapping of Coherence Functions to TIBCO BusinessEvents DataGrid Functions

Table 13 Mapping of Coherence Functions to DataGrid Functions

4.x Coherence Category Function Name	DataGrid Category Function Name
C_CacheGetEntityById()	CacheGetEntityById()
C_CacheLoadConceptByExtId()	CacheLoadConceptByExtId()
C_CacheLoadConceptById()	CacheLoadConceptById()
C_CacheLoadConceptIndexedByExtId()	CacheLoadConceptIndexedByExtId()
C_CacheLoadConceptsByExtId()	CacheLoadConceptsByExtId()
C_CacheLoadEntity()	CacheLoadEntity()
C_CacheLoadEventByExtId()	CacheLoadEventByExtId()
C_CacheLoadEventById()	CacheLoadEventById()
C_CacheLoadParent()	CacheLoadParent()
C_CacheName()	CacheName()
C_CacheReevaluate()	CacheReevaluate()
C_ClassName()	ClassName()
C_EnableCacheUpdate()	EnableCacheUpdate()
C_Index()	Index()
C_Lock()	Lock()
C_TransactionProperties	TransactionProperties()

Table 13 Mapping of Coherence Functions to DataGrid Functions (Cont'd)

4.x Coherence Category Function Name	DataGrid Category Function Name
C_UnLock()	UnLock()

Automatic Refactoring Actions and Limitations

This section explains what is done for each kind of refactoring operation. It also lists some limitations.

Refactoring Limitations

Changes made to certain items are not refactored in this release. You must handle reference updates manually.

References to moved or changed folders in strings References to changed or moved folders are not updated in strings, including CDD and XSLT strings. CDD strings are used in the Cluster Definition Descriptor editor to point to project resources using their project path. XSLT strings are used in mapper functions, which are completed using the Function Argument Mapper. (Such references are, however, updated for entity refactoring operations).

Copy-paste of folders Elements inside the pasted folder are not updated.

References to global variables and shared resources Refactoring does not handle changes to global variables and shared resources. You must manually update references to global variables and shared resources that you change.



TIBCO BusinessEvents destinations and property definitions cannot be moved.

Refactoring for Move and Rename Operations

Move and rename refactoring operations change only the structure of a project. For example, when you change a concept name, that name must change everywhere the concept is referenced in the project. If the element has its own file, the file must also be renamed.

References to the changed or moved element are handled as shown in [Table 14](#). Projects can be complex; this list covers the main cases.

Table 14 Refactoring for Move and Rename Operations

Renaming this...	Updates references in these places...
Concept	<ul style="list-style-type: none"> Concepts that inherit from this one Event expiry actions Property definitions for contained or referenced concepts State models Rules and rule functions
Event	<ul style="list-style-type: none"> Events that inherit from this one Event expiry actions Destination (Default Event) State models Rules and rule functions
Property of a concept or an event	<ul style="list-style-type: none"> Event expiry actions State models Rules and rule functions
Domain Model	<ul style="list-style-type: none"> Domain models that inherit from this one Associated properties
Channel	<ul style="list-style-type: none"> Event default destination paths
Destination	<ul style="list-style-type: none"> Event default destination paths
Rule	<ul style="list-style-type: none"> The rule source Rules and rule functions
Rule Function	<ul style="list-style-type: none"> Event expiry actions State models The rule function source Rules and rule functions

Table 14 Refactoring for Move and Rename Operations (Cont'd)

Renaming this...	Updates references in these places...
Folder	All location that this folder is used in a path, for example path to a default destination in an event, property definitions for contained or referenced concepts, and in rules and rule functions. For folder refactoring limitations, See Refactoring Limitations on page 53 .

Refactoring for Delete Operations

Element deletion can affect project behavior. You must take care to ensure that your project behavior is as desired after the deletion refactoring.

Deletion removes all references to the deleted object only in specific cases:

- References to a deleted domain model are removed from concepts that refer to it.
- Reference to a deleted state model are removed from the owning concept

Channels and Destinations

This chapter explains how to configure channels and destinations, and includes information about Rendezvous and local channels. Additional information about JMS channel is provided in [Chapter 5, JMS Channels, on page 83](#), HTTP channel in [Chapter 6, HTTP and SOAP Channels, on page 97](#), Hawk channels in [Chapter 7, HAWK Channel, on page 141](#), and ActiveSpaces channel in [Chapter 8, ActiveSpaces Channel, on page 148](#).

Topics

- [Overview of Channels and Destinations, page 58](#)
- [Selecting a Serializer, page 60](#)
- [Mapping Incoming Messages to Non-default Events, page 61](#)
- [Working with Rendezvous Channels, page 62](#)
- [Working with Local Channels, page 65](#)
- [Adding Channels and Destinations, page 66](#)
- [Communicating with Other Sources using TCP, page 70](#)
- [Channel Resource Reference, page 73](#)
- [Destination Resource Reference, page 78](#)

Overview of Channels and Destinations

One project can have multiple channels of different types with multiple destinations as needed.

See Also

- Chapter 2, Channels and Events in *TIBCO BusinessEvents Architect's Guide* to understand more about the role of channels in a project, and how messages are acknowledged.
- [Chapter 41, Diagrams, page 701](#) for information on using channel dependency diagrams.
- For TRA file updates required for Rendezvous and JMS channels that use TIBCO Enterprise Message Service, see the following section in *TIBCO BusinessEvents Administration: For TIBCO Enterprise Message Service and TIBCO Rendezvous Channels*.
- For configuring SSL over HTTP when the Identity Resource is of the type Certificate/KeyURL, you need to setup native Tomcat libraries. See the section Adding Installation or Build Script for OpenSSL for Tomcat in the *TIBCO BusinessEvents Installation* guide.

Types of Channels

You can choose from the following types of channels:

- **TIBCO Rendezvous channels** Connect TIBCO BusinessEvents to TIBCO Rendezvous sources and sinks. See [Working with Rendezvous Channels on page 62](#).
- **Local channels** Connect co-located agents at runtime. See [Working with Local Channels on page 65](#).
- **JMS channels** Connect TIBCO BusinessEvents to TIBCO Enterprise Message Service provider sources and sinks. See [Chapter 5, JMS Channels, on page 83](#).
- **HTTP channels, including SOAP support** An HTTP channel acts as an HTTP server at runtime. This enables TIBCO BusinessEvents to serve requests from clients, and to act as a client of other servers. Support for SOAP protocol is provided by a specialized event type and a set of functions for extracting information from SOAP request messages and constructing response messages. See [Chapter 6, HTTP and SOAP Channels, on page 97](#).
- **Hawk channels** Connect TIBCO BusinessEvents to TIBCO Hawk domain. This enables TIBCO BusinessEvents to receive events from the Hawk monitor and

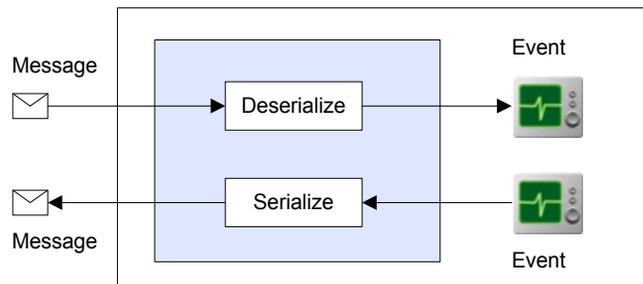
transform them into events. A set of catalog functions are also provided which are used to control the Hawk microagents through Hawk APIs. See [Chapter 7, HAWK Channel, on page 141](#) for details.

- **ActiveSpaces channels** Connect TIBCO BusinessEvents to TIBCO ActiveSpaces metaspaces. This enables TIBCO BusinessEvents to monitor the activities on the TIBCO ActiveSpaces metaspaces, receive events from TIBCO ActiveSpaces and convert them into events in TIBCO BusinessEvents. A set of catalog functions are provided to control the type of the TIBCO ActiveSpaces events. See [Chapter 8, ActiveSpaces Channel, on page 148](#) for details.

Selecting a Serializer

For each type of channel (except local channels), TIBCO BusinessEvents uses a serializer to convert events to messages and a deserializer to convert incoming messages to events. Local channels do not require serializers. HTTP channels also provides you the option of using action rule functions on the message instead of converting messages to event using deserializer.

Figure 1 *Serializer and Deserializer Behavior*



When you configure a destination, you select the appropriate serializer. (It actually includes the serializer and deserializer).

For JMS and Rendezvous messages, different serializers are available for different needs. For details see [Selecting a JMS Serializer on page 85](#) and [Working with Rendezvous Channels on page 62](#).

For HTTP, and SOAP, specific serializers are used. See [Mapping of HTTP Requests to Events, with the REST serializer on page 99](#) for HTTP channel configuration. See [Configuring TIBCO BusinessEvents as a SOAP Server and Client on page 122](#) for more on SOAP.

For TIBCO ActiveSpaces tuples, the AS Serializer is used. See [Working with ActiveSpaces Channels on page 155](#) for ActiveSpaces channel configuration.

For TIBCO Hawk, see [Working with the Hawk Channel on page 142](#) for details.

Deploy-time Configuration

In the Cluster Deployment Descriptor, you specify what destinations an agent listens to at runtime, and you configure the destinations with an event preprocessor (as needed) and threading options. See the section [Configuring Collections of Rules, Rule Functions, and Destinations on page 492](#).

Mapping Incoming Messages to Non-default Events

Incoming messages can be mapped to a default event that you specify when you configure the destination. (See *Default Destinations and Default Events* in *TIBCO BusinessEvents Architect's Guide* for more details). You can also map incoming messages to specified event types.

The fields in a message header instruct TIBCO BusinessEvents to map the incoming message to a specified event type:

- The field named `_ns_` takes a namespace as a value. The namespace points to the event type, for example, `www.tibco.com/be/ontology/Events/MyEvent`
- The field named `_nm_` takes the name of the event, for example, `NewMyEvent`
- The field named `_extid_` takes the unique external id of the event.

These fields are added and filled automatically for events created using TIBCO BusinessEvents rules. You can also add these fields to messages coming in from other sources, if you have control of those messages.

See also:

[Mapping of HTTP Request URI to Destination on page 100](#)

[Mapping of SOAP Request URI to Destination on page 123](#)

[Specifying Default and Non-Default Destinations on page 164](#)

Working with Rendezvous Channels

This section explains the behavior of the provided Rendezvous serializer classes and other related information.

Note that the Rendezvous serializers use UTF8 encoding for XML payloads.

The type of the event from which the outbound message is serialized is added to the Rendezvous message header using `_nm_` and `_ns_` fields so that if the message is used in TIBCO BusinessEvents again, the correct event type is used to deserialize the message (ignoring the default event specified for the destination).

See [Adding Channels and Destinations on page 66](#) for details on adding Rendezvous channels.

For TRA file updates required for Rendezvous and JMS channels that use TIBCO Enterprise Message Service, see the following section In *TIBCO BusinessEvents Administration*: For TIBCO Enterprise Message Service and TIBCO Rendezvous Channels ().

Rendezvous Message Header

For Rendezvous messages, the only header that TIBCO BusinessEvents interprets is `_sendsubject_` which is of type String. It is a read-only property. The event has to define this property to receive the value. The value is the actual Subject on which the message was sent.

Basic Serializer

The basic serializer, `TibRvMsgSerializer` serializer, is for efficient handling of events and messages that do not have payloads. It ignores payloads in messages and in events if any exist.

First level Rendezvous property values are used as values for matching event properties. Any additional (non-matching) Rendezvous properties are ignored.

Serializer For Use with Payloads

To include a payload in a Rendezvous message, ensure that the message has a `_payload_` field. To pass contents between the Rendezvous message `_payload_` field and an event payload, use the `TibRvMsgSerializerWithXMLPayload` serializer.

Deserializing from Rendezvous Message to Event

First level Rendezvous property values are used as values for matching event properties. Any additional (non-matching) Rendezvous properties are ignored.

The `_payload_` field contents are passed into the event payload. Supported `_payload_` field datatypes and Rendezvous wire format types are as follows:

Data Type	Wire Format Type
String	TibrvMsg.STRING
TibrvXML	TibrvMsg.XML
byte[]	TibrvMsg.OPAQUE TibrvMsg.I8ARRAY

If the event defines a payload, but the incoming Rendezvous message does not have a `_payload_` field, TIBCO BusinessEvents attempts to map the entire message as the event payload.

Serializing from Event to Rendezvous Message

Event properties are transformed to first level Rendezvous message properties.

The event payload is passed to the Rendezvous message `_payload_` field.

If the `_payload_` field is of an unsupported type, or is missing, or if the event has not been configured for a payload, the payload is ignored.

Avoiding REGISTRATION COLLISION RVCN Advisory Messages

This section applies to configuring RVCN when using a Rendezvous transport shared resource (see [Rendezvous Transport on page 245](#)).

REGISTRATION.COLLISION RVCN advisory messages can be thrown when when the same CM name (or CMQ name) is used on all engines. These name collisions can result in thrashing.

To prevent this issue, you must ensure that the CM names (and CMQ names) are different on all engines. To do so, add global variables to the CM Name or CMQ Name, and to the ledger file name if a ledger file is used for RVCN, to ensure the uniqueness of these names. Add one or more of the following variables, depending on need (as explained below):

```
%%EngineName%%
%%ChannelName%%
%%ChannelURI%%
```

The `%%EngineName%%` variable is generally required for all names. Note that you must start engines using unique names so that the value of each engine's `%%EngineName%%` variable be different at runtime.

In addition, if different channels use the same RVCN shared resource, you also need to add `%%ChannelName%%` or `%%ChannelURI%%`.

Use `%%ChannelURI%%` in cases where channels using the same RVCN shared resource have the same name but are in different folders.

You must define any of the above String type global variables you use. They are not predefined. However, TIBCO BusinessEvents provides the value at runtime, so you can use any string value or use an empty string as the value when you define the variables.

Working with Local Channels

Local channels are used in rules or rule functions to route events to an appropriate agent running in the same engine (processing unit). Local channels are useful in two cases:

- For applications using In Memory object management (generally used only for testing)
- For certain scenarios where an inference agent is co-deployed with a query agent. See *TIBCO BusinessEvents Event Stream Processing Query Developer's Guide*.

See [Adding Channels and Destinations on page 66](#) for details on adding local channels.

Event Use Count

Local channels pass the same event object between the agents. Consuming the event in one agent does not affect other agents that also received the event over a local channel. A use count is maintained for each event to track how many agents have received the event but not consumed it. The use count of the event is incremented depending on the number of agents it is routed to. When an event is consumed in one agent, TIBCO BusinessEvents deletes the reference to the event in that agent and it decrements the use count of the original event instance.

Using Local Channel Destinations

Local channel destinations can use an event preprocessor.

Local channel destinations do not have default events.

To route an event to a local channel's destination, use the `Event.routeTo()` function. (You can use this function for other purposes too.)

In the provided example, `BE_HOME/Examples/MultipleSessionsAndLocalChannel`, events containing small orders are sent to the agent that deals with small orders as follows:

```
Event.routeTo(order, "/Channels/Local/toSmall", "");
```

The signature of this function is as follows:

```
SimpleEvent routeTo(SimpleEvent event, String destinationPath, String properties)
```

Adding Channels and Destinations

The general procedure for creating channels and destinations of all types is the same, though the configuration options are different. A TIBCO Rendezvous channel is shown as an example.

The screenshot shows the configuration interface for a TIBCO Rendezvous channel. The main window is titled "Channel: RV". It is divided into several sections:

- Configuration:** Includes a "Description" field, a "Driver" dropdown menu set to "RENDEZVOUS", and a "Methods of Configuration" dropdown menu set to "Properties".
- Extended Configuration:** A section that is currently collapsed.
- Properties:** Includes fields for "Service", "Network", and "Daemon".
- Destinations:** This section is expanded to show a list of destinations on the left (containing "input") and a detailed configuration form for the selected destination on the right. The configuration form includes:
 - Name:** input
 - Description:** (empty)
 - Default Event:** /Events/ExternalEvent (with a "Browse..." button)
 - Serializer/Deserializer:** com.tibco.cep.driver.tibrv.serializer.TibRvMsgSerializerWithXMLPayload
 - Subject:** BE.EXAMPLE.EXTERNALEVENT.INPUT
 - RVCM Pre Registration:** (empty)
 - LimitPolicy:** DISCARD_NONE
 - MaxEvents:** 0
 - DiscardAmount:** 0

To Add a Channel

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the channel and select **New > Channel**. You see the New Channel Wizard.
2. In the Channel name field, type a name for the channel. In the Description field, type a description.



You cannot change the name in the editor. To change the name of any project element, right-click the element in Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. In the Driver type field, select the appropriate driver:
 - Rendezvous
 - Local
 - HTTP
 - JMS
 - ActiveSpaces
 - Hawk

If you select Local, channel configuration is complete.
4. Click **Finish**. You see the Channel editor.

To Edit a Channel

1. Edit the Description as needed.
2. In the Driver type field, the driver you selected in the wizard is selected. If you want to change to a different type of channel you can do so now. Select the appropriate driver:
 - Rendezvous
 - Local
 - HTTP
 - JMS
 - ActiveSpaces
 - Hawk
3. For all channel types (except Local and HTTP), from the Method of Configuration drop-down list, select one of the following:
 - **Properties** Select Properties to configure this channel resource using properties. When you select Properties, a Properties section displays appropriate fields for the type of channel. See [Channel Resource Reference on page 73](#) for help in entering the correct values.
 - **Resource** Select Resource if you have a shared resource in your project whose properties you want to reuse for this channel. For HTTP channels, Resource is preselected and cannot be changed.



Resource names and directory names The path to the resource and the resource name cannot contain any of the keywords or other reserved words listed in [Keywords and Other Reserved Words on page 328](#).

4. For HTTP channels only, in the `serverType` field of the Extended Configuration section, TOMCAT is preselected as the server type and cannot be changed.
5. If you selected Resource in the Method of Configuration field, in the Resource field click Browse and select the shared resource you want to use.
6. If you are configuring an HTTP channel, click Advanced and configure tuning properties. See [Configuring TIBCO BusinessEvents to Receive and Send HTTP Requests on page 102](#) for a more detailed procedure and [HTTP Channel Advanced Configuration Settings on page 108](#) for details.
7. Continue to the section [Adding a Destination to a Channel on page 68](#) and add destinations to the channel as needed.

Adding a Destination to a Channel

One channel can have multiple destinations. Each is shown in the Destinations section of the Channel editor.

1. If it is not already open, open the editor for the channel to which you want to add a destination. To open the channel editor, double-click the channel name in the BusinessEvents Studio Explorer view.
2. In the Destinations section, click **Add**.

Common Fields

3. Enter a Name and Description for the destination.
4. In the Default Event field, browse to and select the event to be created (by default) from incoming messages received by this destination. If you have not yet created the event, you can select the default event later.
5. In the Serializer/Deserializer field, select the appropriate class. See the following sections for more details:
 - For JMS destinations: [Selecting a JMS Serializer on page 85](#)
 - For Rendezvous destinations: [Working with Rendezvous Channels on page 62](#)
 - For HTTP and SOAP destinations: [Task D, Add a Destination, on page 104](#), in the section [Configuring TIBCO BusinessEvents to Receive and Send HTTP Requests on page 102](#). (Destination configuration is the same for SOAP and HTTP destinations. See [Manually Creating Resources to Work with SOAP Services on page 124](#) and [Creating Resources Using the WSDL](#)

[Import Utility on page 125](#) for more information about creating SOAP channels and destinations)

- For Hawk destinations: see [Working with the Hawk Channel on page 142](#).
- For ActiveSpaces destinations: [Working with ActiveSpaces Channels on page 155](#).

Various fields appear, depending on your selection.

Channel-Specific
Fields

6. Complete the rest of the properties for the type of destination you are creating. See [Destination Resource Reference on page 78](#) for details.
7. Save the project.

Communicating with Other Sources using TCP

In addition to Channels, TIBCO BusinessEvents can also communicate with other data sources using TCP. You can create a local TCP server and a TCP client so that TIBCO BusinessEvents can communicate with data sources not otherwise available through channels, using TCP.

TCP communication is available as a Communication Built-in Function in the Catalog Functions view. Using this set of functions you do the following to communicate with TCP servers:

- Create a local TCP server in a startup function using `TCP.createLocalServer()`
- Connect to a Remote TCP server as a client using `TCP.connectToRemoteServer()`
- Register the session listener using `TCP.registerSessionListener()`
- Start the local server using `TCP.startLocalServer()`
- Create callback rule functions and register them as callbacks to the TCP listeners. These callback rule functions create events that are sent to an appropriate destination.
`TCP.readIntoPayload(SessionName)`

Example TCP Rule Function to Start a Local TCP Server

Here is a sample rule function to start a local TCP server:

```
void rulefunction RuleFunctions.InitTCPServers {
    attribute {
        validity = ACTION;
    }
    scope {

    }
    body {
        System.debugOut( "Initializing TCP servers" );
        try {
            TCP.createLocalServer("MyTCPServer", "localhost",
                System.getGlobalVariableAsInt("NSN/SocketAdaptor/Port", 8055));
            TCP.registerSessionListener("MyTCPServer",
                "/RuleFunctions/RawCDRCallback");
            TCP.startLocalServer("MyTCPServer");
        } catch (Exception ex) {
            System.debugOut("Exception occurred while initializing TCP server: " +
                ex@message);
        }
    }
}
```

```

        System.debugOut("TCP server initialization done");
    }
}

```

Example TCP Rule Function to Connect to a Remote TCP Server

Here is a sample rule function to connect to a remote TCP server as a client:

```

Events.RemoteMsgResponseEvent rulefunction RuleFunctions.RemoteTCPSender {
    attribute {
        validity = ACTION;
    }
    scope {
        String host;
        int port;
        String message;
    }
    body {
        String tcpNickName = "TCP-" + host + "-" + port + "-"
            + uri + "-" + closure + "-" + System.nanoTime();
        Events.RemoteMsgRequestEvent requestEvent =
            Events.RemoteMsgRequestEvent.RemoteMsgRequestEvent(null, message);
        TCP.connectToRemoteServer(tcpNickName, host, port);
        TCP.write(tcpNickName, requestEvent);
        TCP.endWrite(tcpNickName);
        Events.RemoteMsgResponseEvent responseEvent =
            TCP.readIntoPayloadFully(tcpNickName,
                "/Events/RemoteMsgResponseEvent");
        TCP.disconnectFromRemoteServer(tcpNickName);
        return responseEvent;
    }
}

```

APIs for TCP Communication

The Catalog Functions view lists the following functions (APIs) for TCP communication:

- connectToRemoteServer()
- createLocalServer()
- disconnectFromRemoteServer()
- disconnectLocalSession()
- endRead()
- endWrite()

- `readIntoPayload()`
- `readIntoPayloadFully()`
- `registerSessionListener()`
- `startLocalServer()`
- `stopLocalServer()`
- `write()`
- Advanced
 - `getReaderInputStream()`
 - `readIntoByteArray()`

Documentation for functions is provided in the tooltips you can see when browsing the functions in TIBCO BusinessEvents Studio. You can also see this documentation in the *TIBCO BusinessEvents Functions Reference*, available in the HTML product documentation.

Channel Resource Reference



Channels allow TIBCO BusinessEvents to listen to and send out messages. Channels contain destinations.

You can configure channels of different types, using the appropriate driver. See [Overview of Channels and Destinations on page 58](#).

To configure an HTTP channel resource, select an HTTP connection resource. No other channel resource fields require configuration.



Local channels are in memory; information in a local channel could be lost if the TIBCO BusinessEvents engine fails.

Wizard and Configuration Section

The Wizard and the Configuration section have the following fields.

Field	Global Var?	Description
Name	No	(Shown in the Wizard and then in the editor title only.) The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. See Identifier Naming Requirements on page 323 .
Description	No	Short description of the resource.
Driver	No	Select the driver for the type of channel you are configuring: <ul style="list-style-type: none"> • TIBCO Rendezvous (see Working with Rendezvous Channels on page 62) • Local (see Working with Local Channels on page 65) • HTTP (See Chapter 6, HTTP and SOAP Channels, on page 97) • JMS (see Chapter 5, JMS Channels, on page 83) • ActiveSpaces (see Chapter 8, ActiveSpaces Channel, on page 148) • Hawk (see Chapter 7, HAWK Channel, on page 141)

Field	Global Var?	Description
Method of Configuration	No	<p>Properties Select Properties to configure this channel resource using properties. See:</p> <ul style="list-style-type: none"> • Configuration for TIBCO Rendezvous Channels on page 74 • Configuration for JMS Channels on page 75 <p>Resource Select Resource if you have a shared resource in your project whose properties you want to reuse for this channel.</p> <p>Note: The path to the resource and the resource name cannot contain any of the words listed in Keywords and Other Reserved Words on page 328.</p>
Resource	No	If you choose Resource as the method of configuration, the Resource field appears. Browse to and select the resource you want to use. As a convenience, you can open the selected resource by clicking the underlined label.
Configuration for TIBCO Rendezvous Channels		
Service	Yes	<p>The name of the service or port number through which Rendezvous sends messages. In most cases you can leave this field empty, accepting the default value.</p> <p>For more information about the Rendezvous service parameter, see <i>TIBCO Rendezvous Concepts</i> or <i>TIBCO Rendezvous Administration</i>.</p> <p>Default is 7500 (defined in Global Variables).</p>
Network	Yes	<p>The network over which Rendezvous sends messages. In most cases you can leave this field empty. For more information about the network parameter, see <i>TIBCO Rendezvous Administration</i>.</p> <p>Default is an empty string (defined in Global Variables).</p>
Daemon	Yes	<p>The location of the Rendezvous daemon, which is usually expressed as a client socket number, for example 6555. In most cases, you can leave this field empty, accepting the default value. For more information about the daemon parameter, see <i>TIBCO Rendezvous Concepts</i>.</p> <p>Default is <code>tcp: 7500</code> (defined in Global Variables).</p>

Field	Global Var?	Description
Configuration for JMS Channels		
ProviderURL	Yes	The URL at which TIBCO BusinessEvents can contact the Enterprise Message Service server. Example: <code>tcp://localhost:7222</code>
UserName	Yes	A valid username for the Enterprise Message Service server.
Password	Yes	The password assigned to the username, above, for the purpose of accessing the Enterprise Message Service server.
IsTransacted	Yes	Accepts <code>true</code> or <code>false</code> . Specify <code>true</code> if the session has transaction semantics. Specify <code>false</code> if it has non-transaction semantics. For more information about the <code>IsTransacted</code> property, see TIBCO Enterprise Message Service documentation.
ClientID	Yes	The unique client ID of the connection.
Configuration for ActiveSpaces Channels		
See Basic TIBCO ActiveSpaces Concepts and Terminology on page 149 for the fundamental concepts and terminology.		
Metaspace Name	No	The name of a particular metaspace instance in TIBCO ActiveSpaces that the channel must connect to. The metaspace must be created and initialized before the channel can use it at runtime.

Field	Global Var?	Description
Discovery URL	Yes	<p>Specifies how a metaspaces instance discovers the current metaspaces members. Multicast discovery can use either PGM - Pragmatic General Multicast or RV - TIBCO Rendezvous protocol.</p> <p>If using PGM protocol, the multicast URL is expressed in the following format:</p> <pre>tibpgm://destination port/interface;discovery group address/optional transport arguments, where</pre> <ul style="list-style-type: none"> <i>destination port</i> specifies the destination port used by the PGM transport. If not specified, the default value of 7888 is used. <i>interface;discovery group address</i> specifies the address of the interface to be used for sending discovery packets, and the discovery group address to be used. If not specified, it defaults to the default interface and discovery address, 0.0.0.0;239.8.8.8. <i>optional transport arguments</i> specifies a semicolon-separated list of optional PGM transport arguments. By default, the PGM transport is tuned to provide the best performance according to the most common deployment architectures. The values of those optional arguments should be changed only when necessary and with care, since any inappropriate values could easily result in degraded performance of the product. <p>If using TIBCO Rendezvous, the multicast URL is expressed in the following format:</p> <pre>tibrv://service/network/daemon, where</pre> <ul style="list-style-type: none"> <i>service</i> specifies the TIBCO Rendezvous service number (UDP port) that will be used. If not specified, the value defaults to 7889. <i>network</i> specifies the interface and discovery group address that will be used to send Rendezvous discovery packets. The format is: <code>interface;discovery_group_address</code> If not specified, the default interface and discovery group address will be used. (<code>tibrv://7889/;239.8.8.9/</code>). <i>daemon</i> specifies where to find the Rendezvous daemon. If not specified, it will try to connect to a local daemon on port 7500.

Field	Global Var?	Description
ListenUrl	Yes	<p>The discovery mechanism is based on pure TCP.</p> <p>All the designated well known metaspace members are identified by an IP address and a port number. This address and port are specified by the member's Listen URL.</p> <p>If not specified, the discovery process uses the default IP address and the first free TCP port that can be acquired from the operating system (starting 5000 and above).</p>

Destination Resource Reference



Within each channel, destinations direct incoming and outgoing information. A channel resource is not ready to use until it has at least one destination.

The Destinations section of a channel has the following fields.

Field	Global Var?	Description
Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. See Identifier Naming Requirements on page 323 .
Description	No	Short description of the resource.
Default Event	No	The event to be created from incoming messages unless otherwise specified. As a convenience, you can open the selected event resource by clicking the underlined label. Optional, but only if you always specify an event type in the incoming message. Not used for local channel.
Serializer/ Deserializer	No	Specify a serializer class to convert messages to simple events and simple events to messages. See the following sections for details: <ul style="list-style-type: none"> • Selecting a JMS Serializer on page 85 • Working with Rendezvous Channels on page 62 • Task D, Add a Destination, on page 104
Configuration for TIBCO Rendezvous Destinations		
See TIBCO Rendezvous documentation for more details on these settings.		
Subject	Yes	The TIBCO Rendezvous subject for incoming and outgoing messages.

Field	Global Var?	Description (Cont'd)
RVCM Pre Registration	Yes	For TIBCO Rendezvous certified message publishers, specify pre-registered listener names as a comma separated list.
LimitPolicy	Yes	How you want the Rendezvous listener to behave when it receives more messages than the MaxEvents limit. Choose one of: Discard_None (default) Discard_First Discard_Last Discard_New When MaxEvents or DiscardAmount are zero (unlimited), the LimitPolicy must be Discard_None.
MaxEvents	No	Maximum number of message events that the queue can hold. The default value, zero (0) means an unlimited number of events.
DiscardAmount	No	The number of events to discard when the queue exceeds its maximum event limit. The default value, zero (0) means events are never discarded.
Configuration for JMS Destinations		
See TIBCO Enterprise Message Service documentation for more detail on these settings.		
Queue	No	Specifies whether the destination is a queue or a topic. Check the checkbox if the destination is a queue. If the destination is a topic, do not check it.

Field	Global Var?	Description (Cont'd)
Name	Yes	<p>Required. The name of the queue or topic.</p> <p>(TIBCO BusinessEvents ignores JMS destinations with null or empty-string queue or topic names. It logs an error message for the ignored destinations. If a JMS message is sent out through an ignored destination, TIBCO BusinessEvents throws an exception and the message is not sent out. TIBCO BusinessEvents also does not receive JMS messages (events) through these ignored destinations.)</p>
Selector	Yes	<p>Specifies a filter to pick up messages from the destination. This is a standard JMS selector based on SQL92 semantics.</p>
DeliveryMode	No	<p>The delivery mode property instructs the server concerning persistent storage for the message. Select one of the following:</p> <p>PERSISTENT (default) — In JMS message headers this is represented by the code 2.</p> <p>NON-PERSISTENT — In JMS message headers this is represented by the code 1.</p> <p>RELIABLE — This value is an extension to the standard, used in TIBCO Enterprise Message Service. In message headers this is represented by the code 22.</p> <p>You can also set a delivery mode in an event. See Using JMS Header Properties in Incoming and Outgoing Messages on page 93.</p>
AckMode	No	<p>The acknowledgement mode. See Setting the JMS Message Acknowledgement Mode on page 89 for a table explaining the various modes.</p> <p>You can also set an acknowledgement mode in a node's engine properties. The setting in the destination overrides the engine property setting.</p> <p>Default is EXPLICIT_CLIENT_ACKNOWLEDGE</p>

Field	Global Var?	Description (Cont'd)
Priority	No	<p>The message priority. Takes a numerical value between 0 and 9. Larger numbers represent higher priority.</p> <p>You can also set a priority in an event. See Using JMS Header Properties in Incoming and Outgoing Messages on page 93.</p> <p>Default is 4</p>
TTL	No	<p>The length of time that the message will live (in milliseconds) before expiration. If set to 0, the message does not expire.</p> <p>You can also set a TTL (<code>JMSExpiration</code>) in an event. See Using JMS Header Properties in Incoming and Outgoing Messages on page 93.</p> <p>Default is 0</p>
Durable Subscriber Name	Yes	<p>For destinations that are JMS Topics, if you provide a <code>DurableSubscriberName</code>, the destination becomes a JMS durable topic subscriber with the specified name. If you do not provide a value, the destination becomes a non-durable topic subscriber.</p> <p>The value of this property can be any unique string and can include any global variables. TIBCO BusinessEvents provides a set of case-sensitive variables that produce a unique <code>DurableSubscriberName</code> string. See Creating Unique JMS DurableSubscriber Name Properties on page 87 for details.</p> <p>Default is:</p> <pre>%%EngineName%%:%%SessionName%%:%%ChannelURI%%:%%DestinationName%%</pre>

Configuration for Local Destinations

Local destinations do not use serializers, deserializers, or default events.

Size	No	<p>The maximum number of events to be held in the queue. The default is zero (0), which allows unlimited events in the queue.</p>
------	----	---

Field	Global Var?	Description (Cont'd)
TimeOut	No	Time to wait when sending an event to this local destination: -1 Waits indefinitely 0 Does not wait >0 Waits for the number of milliseconds specified Default is -1.

Configuration for HTTP Destinations

Is Page Flow	No	Enables Action Rule Function based approach. If checked, the system disables the Default Event and Serializer/Deserializer fields for input.
ContextPath	No	Context URI for the web application. The field is active for input if the Is Page Flow checkbox is checked.
Action Rule Function	No	The Action Rule function to be executed when an HTTP message arrives at the context URI for the web application. The field is active for input if the Is Page Flow checkbox is checked.

Chapter 5 **JMS Channels**

See [Chapter 4, Channels and Destinations, on page 57](#) for basic information and procedures that apply to all types of channels. This chapter provides additional information about working with JMS channels.

Topics

- [Overview of JMS Channels, page 84](#)
- [Selecting a JMS Serializer, page 85](#)
- [Creating Unique JMS DurableSubscriber Name Properties, page 87](#)
- [Setting the JMS Message Acknowledgement Mode, page 89](#)
- [When JMS Messages are Acknowledged, page 92](#)
- [Using JMS Header Properties in Incoming and Outgoing Messages, page 93](#)
- [JMS Header Field Names, page 95](#)

Overview of JMS Channels

This chapter addresses tasks you may have to complete when configuring a JMS channel.

- To understand how to choose a serializer to handle the JMS message types that will be sent to the destination, see [Selecting a JMS Serializer on page 85](#)
- To understand how you can create unique DurableSubscriber Name property values for JMS Topic destinations, see [Creating Unique JMS DurableSubscriber Name Properties on page 87](#).
- To learn how to change the JMS message acknowledgement type from the default type, see [Setting the JMS Message Acknowledgement Mode on page 89](#).
- To understand when JMS messages are acknowledged, see [When JMS Messages are Acknowledged on page 92](#)
- To understand how TIBCO BusinessEvents handles certain JMS header properties, and how you can work with these header properties, see [Using JMS Header Properties in Incoming and Outgoing Messages, page 93](#) and the related reference section, [JMS Header Field Names on page 95](#).

See Also

- [CDD Processing Units Tab JMS Server Connection Properties on page 520](#) for details about run-time properties used to reconnect to a JMS server.
- For TRA file updates required for Rendezvous and JMS channels that use TIBCO Enterprise Message Service, see the following section In *TIBCO BusinessEvents Administration*: For TIBCO Enterprise Message Service and TIBCO Rendezvous Channels ().

Selecting a JMS Serializer

This section explains the purpose of the provided JMS serializer classes.

Choose the serializer that handles the JMS message types that will be sent to the destination you are configuring.

Table 15 Which Serializers to Use for JMS Message Types

Message Type	BytesMessage Serializer/ UtfBytesMessage Serializer	TextMessage Serializer	MessageWithNoBody ¹
Message	Supported	Supported	Supported
BytesMessage	Supported	(N/A)	Supported
MapMessage	Supported	Supported	Supported
ObjectMessage	Not Supported	Not Supported	Not Supported
StreamMessage	Not Supported	Not Supported	Not Supported
TextMessage	(N/A)	Supported	Supported

1. MessageWithNoBody corresponds to the BasicMessageSerializer option in TIBCO BusinessEvents Studio.

For MapMessage messages, you create properties whose names match the message keys.

All serializers support reading and writing application header properties and JMS header properties. The difference between the serializers is in how they handle payloads.

BytesMessageSerializer decodes the body as a sequence of bytes and parses them to create an XML structure according to the payload definition in the event.

UtfBytesMessageSerializer is similar to the BytesMessageSerializer except that it serializes the payload using writeUTF() instead of writeBytes(), and deserializes the payload using readUTF() instead of readBytes().

TextMessageSerializer decodes the text from the message as an XML string.

MessageWithNoBody serializer does not serialize or deserialize the payload.

The BytesMessageSerializer is used to receive JMS messages coming in as a stream of uninterrupted bytes in the message body.

Typically, `BytesMessageSerializer` is a low level serializer and can have performance implications. TIBCO recommends that you choose the serializer for a defined message type (`MapMessage`, `TextMessage`, and so on) that most closely matches the expected usage.

See [Using JMS Header Properties in Incoming and Outgoing Messages on page 93](#).

BytesMessageSerializer

For incoming messages of type JMS `BytesMessage`, the serializer converts the message bodies to event payloads. The payloads are XML String type, but are not human-readable. However, you can access them using XPath functions.

For outgoing events, the serializer converts XML payloads to JMS `BytesMessage` message bodies.

The `BytesMessageSerializer` class is the default serializer.

UtfBytesMessageSerializer

The `UtfBytesMessageSerializer` is similar to the [BytesMessageSerializer](#), except that it serializes the payload using `writeUTF()` and deserializes the payload using `readUTF()`.

For outgoing events, the serializer converts XML payloads to JMS `BytesMessage` message bodies.

TextMessageSerializer

For incoming messages, the `TextMessageSerializer` serializer converts JMS `TextMessage` messages to event payloads. The payloads are XML String type, and are human-readable. You can access them using XPath functions.

For outgoing events, the serializer converts XML payloads to JMS `TextMessage` messages.

MessageWithNoBody

The `MessageWithNoBody` serializer does not serialize or deserialize the payload. This serializer corresponds to the `BasicMessageSerializer` option in TIBCO BusinessEvents Studio.

For outgoing events, the serializer converts the payload to messages of type `Message`.

Creating Unique JMS DurableSubscriber Name Properties

For destinations that are JMS Topics, if you provide a DurableSubscriber Name when you configure the destination resource, the destination becomes a JMS durable topic subscriber with the specified name. This section explains how you can ensure that the DurableSubscriber Name value is unique.



- When using topic destination with a durable name in applications using In Memory OM and fault tolerance, do not provide a value for the Client ID setting and do not check the Auto-generate Client ID check box in the JMS shared resource.
- Do not use durable topic destinations for multi-agent applications, even when only one agent instance is active at a time (that is, even when Agent Classes tab > *AgentClassName* > Max Active is set to 1). Instead, use queue destinations.

The value of the DurableSubscriber Name property can be any unique string and can include any global variables. TIBCO BusinessEvents provides a set of case-variables that produce a unique DurableSubscriberName string:

```
%%Deployment%%:%%EngineName%%:%%SessionName%%:%%ChannelURI%%:%%DestinationName%%
```

The first variable, %%Deployment%%, is a standard TIBCO global variable. The other three are only for use with the DurableSubscriberName property within TIBCO BusinessEvents. See [Table 16, Variables for Use with DurableSubscriberName](#), for details.



Do not attempt to use %%EngineName%%, %%SessionName%%, %%ChannelURI%%, or %%DestinationURI%% in any area of TIBCO BusinessEvents software except the DurableSubscriberName property.

Table 16 Variables for Use with DurableSubscriberName

Variable	Description
%%EngineName%%	The name of the TIBCO BusinessEvents engine. The name used is established using a series of checks. See <i>Determining the Engine Name</i> in <i>TIBCO BusinessEvents Administration</i> for details.
%%SessionName%%	The name of the agent class that is associated with the durable subscriber. Agent classes are defined in the CDD resource. See Chapter 29, Agent and Processing Unit Configuration, on page 489 for details.

Table 16 Variables for Use with `DurableSubscriberName` (Cont'd)

Variable	Description
<code>%%ChannelURI%%</code>	The path to the channel within the TIBCO BusinessEvents project: <i>/folder/.../channel_name</i>
<code>%%DestinationName%%</code>	The name of the TIBCO BusinessEvents destination, within the channel specified in <code>%%ChannelURI%%</code> .

Setting the JMS Message Acknowledgement Mode

JMS channels support connection to TIBCO Enterprise Message Service destinations. The default acknowledgement mode is EXPLICIT.

You can set the acknowledgement mode either of these two ways:

- In the AckMode field of a destination resource.
- Using a configuration property.

The destination setting overrides the property setting. To set the acknowledgement mode using a configuration property, use one of the following, depending on whether you are using topics or queues:

```
be.channel.tibjms.topic.ack.mode acknowledgement_mode_number
```

```
be.channel.tibjms.queue.ack.mode acknowledgement_mode_number
```

Where *acknowledgement_mode_number* is one of the numbers that represent an acknowledgement mode, as shown in [Table 17](#).

Note that in TIBCO Enterprise Message Service, mode names are slightly different. They are prefixed with TIBEMS-. See *TIBCO Enterprise Message Service User's Guide* for more details.

Table 17 JMS Message Acknowledgement Modes

No.	Mode	Description
1	AUTO_ACKNOWLEDGE	Specifies that the session is to automatically acknowledge consumer receipt of messages when message processing has finished.
2	CLIENT_ACKNOWLEDGE	Specifies that the consumer is to acknowledge all messages that have been delivered so far by the session. When using this mode, it is possible for a consumer to fall behind in its message processing and build up a large number of unacknowledged messages. See Using CLIENT_ACKNOWLEDGE Mode with Websphere MQ and Cache-Aside on page 91 for required configuration for Websphere MQ when cache-aside database write strategy is used.

Table 17 JMS Message Acknowledgement Modes (Cont'd)

No.	Mode	Description
3	DUPS_OK_ACKNOWLEDGE	<p>Specifies that the session is to lazily acknowledge the delivery of messages to the consumer. "Lazy" means that the consumer can delay acknowledgement of messages to the server until a convenient time; meanwhile the server might redeliver messages. This mode reduces session overhead. However, should JMS fail, the consumer may receive duplicate messages.</p>
23	EXPLICIT_CLIENT_ACKNOWLEDGE (TIBCO Proprietary)	<p>TIBCO Enterprise Message Service extension to JMS acknowledge modes.</p> <p>This is the default.</p> <p>EXPLICIT_CLIENT_ACKNOWLEDGE is like CLIENT_ACKNOWLEDGE except it acknowledges only the individual message, rather than all messages received so far on the session.</p> <p>One example of when EXPLICIT_CLIENT_ACKNOWLEDGE would be used is when receiving messages and putting the information in a database. If the database insert operation is slow, you may want to use multiple application threads all doing simultaneous inserts. As each thread finishes its insert, it can use EXPLICIT_CLIENT_ACKNOWLEDGE to acknowledge only the message that it is currently working on.</p>
24	EXPLICIT_CLIENT_DUPS_OK_ACKNOWLEDGE (TIBCO Proprietary)	<p>TIBCO Enterprise Message Service extension to JMS acknowledge modes.</p> <p>EXPLICIT_CLIENT_DUPS_OK_ACKNOWLEDGE mode is like TIBEMS-DUPS-OK-ACKNOWLEDGE except it "lazily" acknowledges only the individual message, rather than all messages received so far on the session.</p>

Table 17 JMS Message Acknowledgement Modes (Cont'd)

No.	Mode	Description
22	NO_ACKNOWLEDGE (TIBCO Proprietary)	<p>TIBCO Enterprise Message Service extension to JMS acknowledge modes.</p> <p>Suppresses the acknowledgement of received messages. After the server sends a message to the client, all information regarding that message for that consumer is eliminated from the server. Therefore, there is no need for the client application to send an acknowledgement to the server about the received message. Not sending acknowledgements decreases the message traffic and saves time for the receiver, therefore allowing better utilization of system resources.</p> <p>Note Sessions created in NO_ACKNOWLEDGE receipt mode cannot be used to create durable subscribers.</p> <p>Note Also, queue receivers on a queue that is routed from another server are not permitted to specify NO_ACKNOWLEDGE mode.</p>

Using CLIENT_ACKNOWLEDGE Mode with Websphere MQ and Cache-Aside

The cache-aside database write strategy is multi-threaded. However, when Websphere MQ messages are sent using CLIENT_ACKNOWLEDGE_MODE, each message must be handled from start to finish using a single thread. To address this issue do the following:

1. Define the destination using Caller's Thread in the Threading Model setting (in the CDD Collections tab or Agent Classes tab).
See [CDD Collections Tab Input Destination Settings Reference on page 494](#).
2. To ensure sequential operations set the following property in the CDD file at the appropriate level:

```
Agent.agentClassName.enableParallelOps=false
```

Setting this property to false means that all post-RTC operations are done on a single thread.

When JMS Messages are Acknowledged

When TIBCO BusinessEvents acknowledges JMS messages depends on the JMS acknowledgement mode, time to live (TTL) setting, and object management (OM) type, as shown in the table below.

Table 18 When JMS Messages are Acknowledged

JMS Acknowledgement Mode	OM Type	Acknowledged
AUTO_ACKNOWLEDGE DUPS_OK_ACKNOWLEDGE	—	On receipt
NO_ACKNOWLEDGE	—	Never
CLIENT_ACKNOWLEDGE EXPLICIT_CLIENT_ACKNOWLEDGE EXPLICIT_CLIENT_DUPS_OK_ACKNOWLEDGE	Cache	Post RTC
	In Memory	¹ When retracted

1. That is, when the event is deleted (using the `Event.consumeEvent` function) or when the event reaches the end of its time to live (TTL) period.

For more details about message acknowledgment with Cache OM, see Post-RTC and Epilog Handling and Tuning Options in *TIBCO BusinessEvents Architect's Guide*.

Using JMS Header Properties in Incoming and Outgoing Messages

Information in this section assumes you are familiar with JMS and its header properties. Consult your JMS provider documentation for information. This section explains only how TIBCO BusinessEvents supports use of these properties.

Setting Certain Header Properties in Destinations

In the JMS destination Configuration section, you configure the following three header properties:

- DeliveryMode (`JMSDeliveryMode`)
- Priority (`JMSPriority`)
- TTL (`JMSExpiration`)



JMS header properties defined in events take precedence over properties defined in destinations.

Setting Header Properties Using Header Properties from Incoming JMS Messages

You can configure events created from incoming JMS messages to have properties that match the JMS header properties. You can then use those event properties to set JMS header properties in outgoing messages.

These event properties must match the JMS header fields. You must use the names as shown in [Table 19, JMS Header Field Names, on page 95](#). You only have to configure event properties for those fields that you want to use. Incoming JMS message header properties will then populate the corresponding TIBCO BusinessEvents event properties.

Setting JMS Header Properties in Outgoing JMS Messages Using Event Properties

Similarly outgoing JMS message header properties will be populated by the corresponding TIBCO BusinessEvents event properties.

Note that the `JMSMessageID` and `JMSTimeStamp` properties are generated when the message is sent. You cannot set these properties manually.

See [How TIBCO BusinessEvents Sets the JMSReplyTo Header in Outgoing Messages on page 94](#) for special handling of the `JMSReplyTo` header.

JMS header properties defined in events take precedence over properties defined in destinations.



You can add the JMS properties to the Base event in your project so that the properties are inherited by all other events.

See [Table 19, JMS Header Field Names, on page 95](#) for details on all properties.

How TIBCO BusinessEvents Sets the JMSReplyTo Header in Outgoing Messages



TIBCO BusinessEvents cannot act as a client in a JMS request-response scenario because TIBCO BusinessEvents currently cannot dynamically create a destination to listen for JMS messages.

If an event has a string type property named `JMSReplyTo` (case sensitive), TIBCO BusinessEvents reads this event property value as a JMS queue or topic name, according to the event's default destination type. TIBCO BusinessEvents looks up the `javax.jms.Destination` on the connected JMS server using this queue or topic name. If TIBCO BusinessEvents cannot find one, it creates a new `javax.jms.Destination` using the given queue or topic name. TIBCO BusinessEvents then sets the `JMSReplyTo` header property of the outgoing JMS message using this destination



If you use the catalog function `Event.replyEvent(requestEvent, replyEvent)` during the RTC in which the *requestevent* is received, then the *replyevent* is sent to the destination in the `JMSReplyTo` header property of the JMS message associated with the *requestevent*.

JMS Header Field Names

The table below shows the names you must use to define event properties corresponding to JMS header field names, as well as some details about the purpose of each property. The property names are not case sensitive.

Table 19 JMS Header Field Names

Field name	Type	Description
JMSDestination	Object	The destination (queue or topic) to which the message is sent.
JMSDeliveryMode	Integer	<p>The delivery mode specified by the sender. This property instructs the server concerning persistent storage for the message. Value can be:</p> <p>2 — interpreted as PERSISTENT (default).</p> <p>1 — interpreted as NON-PERSISTENT.</p> <p>22 — interpreted as RELIABLE. This value is an extension to the standard used in TIBCO Enterprise Message Service.</p> <p>The integer values are interpreted as the text names of delivery modes.</p> <p>You can also set a delivery mode for a destination. See Setting Certain Header Properties in Destinations on page 93.</p>
JMSExpiration	Long	<p>The length of time that the message will live (in milliseconds) before expiration. If set to 0, the message does not expire.</p> <p>You can also set an expiration (TTL) for a destination. See Setting Certain Header Properties in Destinations on page 93.</p>
JMSPriority	Integer	<p>The message priority, a numerical value between 0 and 9. Larger numbers represent higher priority.</p> <p>You can also set a priority for a destination. See Setting Certain Header Properties in Destinations on page 93.</p>

Table 19 JMS Header Field Names (Cont'd)

Field name	Type	Description
JMSMessageID	String	An ID that uniquely identifies each message sent by a provider. A generated value overrides any value set in the corresponding event property.
JMSTimestamp	Long	The time when the message was handed off to a provider to be sent. The message may be sent later than this timestamp value. A generated value overrides any value set in the corresponding event property.
JMSCorrelationID	String	A correlation ID that can be used to link messages. For example, you can link a response message to a request message. Optional.
JMSReplyTo	String	Name of the JMS destination (queue or topic) to send the message reply to. If null, TIBCO BusinessEvents does not set the outgoing message's property. Optional. Note: Do not set the value to an empty string (""). If you do, TIBCO BusinessEvents sets the queue or topic name to an empty string which creates an exception, and the message is not sent. See How TIBCO BusinessEvents Sets the JMSReplyTo Header in Outgoing Messages on page 94.
JMSType	String	The message type identifier, if used by the provider.
JMSRedelivered	Boolean	If this field is set, it is possible that the message was delivered to the client earlier, but not acknowledged at that time.

See [Chapter 4, Channels and Destinations, on page 57](#) for basic information and procedures that apply to all types of channels. This chapter provides additional information about working with HTTP channels, including those configured to serve SOAP messages.

Topics

- [Overview of HTTP and SOAP Channels, page 98](#)
- [Working with HTTP Requests, page 99](#)
- [Setting Up Fault Tolerance for the HTTP Channel, page 106](#)
- [HTTP Channel Advanced Configuration Settings, page 108](#)
- [Using HTTP Functions to Configure HTTP Request Messages, page 116](#)
- [Configuring TIBCO BusinessEvents as a SOAP Server and Client, page 122](#)
- [Parsing and Building SOAP Messages, page 130](#)
- [Understanding the WSDL to Project Resource Mapping, page 135](#)

Overview of HTTP and SOAP Channels

An HTTP channel is an internal HTTP server. When the TIBCO BusinessEvents engine starts, it starts the internal HTTP server, which listens to the requests on the port specified in the HTTP Connection resource.

SOAP Support

SOAP version 1.1 is supported. TIBCO BusinessEvents supports only document/literal type of encoding of SOAP requests.

A SOAP event is an extension of a SimpleEvent. To create a SOAP event, you create a SimpleEvent that inherits from a SOAPEvent. This creates a default schema in the event payload. Then you edit the schema and introduce header and body elements as necessary.

Using an HTTP channel and a destination configured to use the SOAP serializer and deserializer, TIBCO BusinessEvents can act as a web services platform that sends and receives SOAP requests, and performs whatever operations are provided by the web service.

TIBCO BusinessEvents can import a WSDL file and create the required project artifacts based on it, such as events, rules, rule functions, channels, and destinations. For more details, see [Creating Resources Using the WSDL Import Utility on page 125](#).

TIBCO BusinessEvents can also export a SOAP based rule function to a WSDL. The export utility scans the project for rule functions that take a SOAP event as the input, and generates a WSDL operation for each one. For more details, see [Exporting Resources as a WSDL File on page 127](#).

See [Configuring TIBCO BusinessEvents as a SOAP Server and Client on page 122](#) for more details.

Working with HTTP Requests

HTTP requests are parsed and executed using either of the following approaches:

- **Event based approach** HTTP request is mapped to an event using deserializer.
- **Action Rule Function based approach** HTTP request parameters and data are retrieved using HTTP catalog functions and processed using rule function.

The **Is Page Flow** parameter of the destination identifies the approach followed by destination for processing HTTP requests. If the **Is Page Flow** parameter is set to true the HTTP request is processed using the Action Rule Function based approach otherwise Event based approach is used. This section explains the two different approaches for working with HTTP requests, and how to add an HTTP channel and destination.

Event Based Approach

The following paragraphs explain how an HTTP request is mapped to an event. It assumes basic familiarity with HTTP and with the procedure of setting up a channel. See [Adding Channels and Destinations on page 66](#) if you need basic instructions.

Mapping of HTTP Requests to Events, with the REST serializer

`RESTMessageSerializer`, which is set while configuring the channel, maps HTTP requests to TIBCO BusinessEvents. HTTP headers and HTTP parameters in the GET method are mapped to similarly named event properties. When both parameters and headers are specified, parameters take precedence.

Each HTTP Header consists of a name and field value, which are mapped into an event property name and value.

The POST data in the request must match with the payload as defined in the corresponding event. If there is no payload defined for the event, the POST data is translated into a `ByteArray` payload, and is not accessible via the mapper. In other words, only XML payloads are visible via the mapper.

The transfer encoding (`charset`) in the `Content-Type` header indicates what type of transformation is applied to the POST data (message body) to safely transfer content between sender and recipient. If the `Content-Type` header is missing, UTF-8 is used as the default transfer encoding.

When you want the REST serializer to deserialize a GET request into an event with a payload, include the `_payload_` request parameter. The string value of the `_payload_` parameter will always be used as payload in the event.



Avoid sending long requests using GET. For large payload data requests, it is recommended to use the POST method.

For HTTP responses, all event properties are translated to similarly named HTTP headers and the payload is sent as HTTP content.

Sending Non-ASCII Content to Event Properties

The HTTP 1.1 specification states that only ASCII characters can be sent in HTTP headers.

To send non-ASCII event properties in GET methods, use HTTP parameters. HTTP parameters are passed as the `QueryString` of the request URI, that is, the part of the URI that contains data to be passed to web applications.

To decode the `QueryString`, use either the URI Encoding or the body encoding. The body encoding is specified in the `contentType` HTTP header. Select either the URI Encoding or Use Body Encoding for URI setting in the HTTP channel Advanced tab.

Mapping of HTTP Request URI to Destination

An HTTP request URI must map to a valid TIBCO BusinessEvents destination. If not, an error is returned, and the message is discarded.

In the event based approach, once the destination is established, the HTTP message is converted either into an event based upon `_ns_/_nm_`, or into the default event associated with this destination.

In this case TIBCO BusinessEvents server looks for the destination having the same URI as the `requestURI`.

For example, the `requestURI` for the request `https://localhost:7000/Transport/Channel/StudentDestination` is `/Transport/Channel/StudentDestination`. TIBCO BusinessEvents engine maps the request with a destination having URI `/Transport/Channel/StudentDestination` if it exists.

Action Rule Function Based Approach

The following paragraphs explain how a HTTP request is processed in TIBCO BusinessEvents Studio without mapping to an event.

HTTP Request Processing

The TIBCO BusinessEvents server parses the HTTP request received from client. The server extracts the request header and identifies the request URI of the request. The request URI should map to the context specified in the destination defined for the channel. In this case TIBCO BusinessEvents server looks for the destination having the same context path as the requestURI.

For example, the requestURI for the request `https://localhost:7000/Transport/Channel/StudentDestination` is `/Transport/Channel/StudentDestination`. TIBCO BusinessEvents engine maps the request with a destination having context path `/Transport/Channel/StudentDestination` if it exists.

Action rule function defined in the destination uses catalog functions to get required data and parameters from the HTTP request. The rule function processes these parameters and creates a response, similar to a preprocessor. TIBCO BusinessEvents supports all HTTP methods stated in the HTTP 1.1 specification for Action Rule Function based approach.

Sample Code for Action Rule Function

```
void rulefunction RuleFunctions.Callback {
    attribute {
        validity = ACTION;
    }
    scope {
        Object asyncContextObject;
    }
    body {
        //getting servlet request and response objects
        Object servletRequest = HTTP.Servlet.getServletRequest(asyncContextObject);
        Object servletResponse = HTTP.Servlet.getServletResponse(asyncContextObject);
        System.debugOut("##Servlet request method : " +
HTTP.Servlet.Request.getMethod(servletRequest));
        System.debugOut("##Servlet request content : " +
HTTP.Servlet.Request.getRequestContent(servletRequest));
        System.debugOut("##Servlet request Requester Address : " +
HTTP.Servlet.Request.getRequestorAddress(servletRequest));
        System.debugOut("##Servlet request Request URI : " +
HTTP.Servlet.Request.getRequestURI(servletRequest));
        //getting parameters
        String[] params = HTTP.Servlet.Request.getRequestParameters(servletRequest);
        for(int i=0;i<params@length;i++)
        {
            System.debugOut("## Servlet request Parameters : " +
HTTP.Servlet.Request.getRequestParameter(servletRequest,params[i]));
        }
        //getting headers
    }
}
```

```

        System.debugOut("## Servlet request Header Accept : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Accept"));
        System.debugOut("## Servlet request Header Accept-Encoding : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Accept-Encoding"));
        System.debugOut("## Servlet request Header Accept-Language : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Accept-Language"));
        System.debugOut("## Servlet request Header Accept-Charset : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Accept-Charset"));
        System.debugOut("## Servlet request Header Connection : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Connection"));
        System.debugOut("## Servlet request Header User-Agent: " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"User-Agent"));
        System.debugOut("## Servlet request Header Content-Length : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Content-Length"));
        System.debugOut("## Servlet request Header Content-Type : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Content-Type"));
        System.debugOut("## Servlet request Header Host : " +
HTTP.Servlet.Request.getRequestHeader(servletRequest,"Host"));
        HTTP.Servlet.Response.setResponseHeader(servletResponse, "Content-type",
"text/plain");
        HTTP.Servlet.Response.setResponseContent(asyncContextObject, "response",
true);
    }
}

```

Deploying Multiple Web Applications

You can deploy multiple web applications on single channel. Specify the *.WAR file or a valid J2EE web application folder under the Web Applications section in the advanced tab. Specify the context URI and resource path for the web application. Resource path identifies the actual path of web application resources. In case of a *.WAR file resource path is the location of the war file, and in case of the J2EE web application folder the resource path is the location of the base folder of the web application.

Configuring TIBCO BusinessEvents to Receive and Send HTTP Requests

To configure TIBCO BusinessEvents to receive and send HTTP requests, you need to perform the following tasks:

Task A Add an HTTP Connection

Add an HTTP Connection resource to your project. In the Host and Port fields, specify the host and port to which HTTP clients send requests.

In the Host field, enter the name or IP address of the machine running TIBCO BusinessEvents. This is the HTTP server.

In the Port field, enter any available port on the host machine. This is the port on which the server listens for HTTP requests.

To configure an HTTPS (Secure) Connection Check the **SSL** checkbox, click the **Configure SSL** button, and complete the pop-up dialog settings. The server must authenticate to the client. In the Identity field, provide the location of the Server Identity File. (See [Task B](#)).



In one-way SSL, the server authenticates itself to the client using a Server Identity File, but the client does not have to authenticate to the server.

In two-way SSL, the server authenticates itself to the client using a Server Identity File and the client also authenticates to the server using the Client Identity file.

Check the **Requires Client Authentication** checkbox. This enables the **Trusted Certificates Folder** field, where you provide the Client certificates.

Task B For SSL Only — Add an Identity Resource

TIBCO BusinessEvents supports use of an identity file for SSL. Before you configure the secure HTTP connection, add an Identity resource to your project and configure it.

In the URL field, specify the project path of the keystore file, which must be within the project folders.

In the File Type field, specify the keystore file type, and in the Password field, provide the password for the keystore file.

Task C Add an HTTP Channel

Add a channel to your project and configure it as follows:

1. At the New Channel Wizard, provide a name and description, and in the Driver field select HTTP. Click **Finish**.
2. In the Channel editor Channel tab, update the description as desired. The Driver field is set to HTTP (as set in the wizard). The Method of Configuration is preset to Resource and cannot be changed.
3. In the Resource field, browse to the HTTP connection resource you configured in [Task A](#).
4. Click **Advanced**, and configure run-time configuration properties. These properties provide information such as the location of the document root folder. Some properties are for SSL configuration of HTTP Component servers. See [HTTP Channel Advanced Configuration Settings on page 108](#).

Task D Add a Destination

Add a destination to the channel in the usual way. (See [Adding Channels and Destinations on page 66](#) for details.) The fields for destination differs based on different approaches.

To follow the action rule function based approach Check the Is Page Flow checkbox and specify the appropriate **Context Path** and **Action Rule function** for the destination. Specify the context path in the same format as the server would receive the `requestURI` in the HTTP request.

To follow the event based approach In the **Serializer** field select the appropriate serializer:

```
com.tibco.cep.driver.http.serializer.RESTMessageSerializer
```

Specify the default event in the usual way as needed by your project requirements.



HTTP clients of the TIBCO BusinessEvents server would use the complete destination URI, after the host and port for example,
`http://www.acme.com:5560/Transport/Channel/MyDestination.`

See [Mapping of HTTP Request URI to Destination on page 100](#).

Task E Create Events as Needed and Set Default Destinations

Skip this step for action rule function based approach.

In the event based approach For receiving HTTP requests and sending responses, configure events in the usual way, and select an HTTP-based destination as the default destination.

Task F Configure Rules and Rule Functions

Configure rules and rule functions according to your needs and HTTP request processing approach.

In the action rule function based approach For example, in response to a POST request you might do the following:

- Create a servlet request object for the HTTP request using the catalog functions to extract data from the HTTP request. See [Chapter 18, Functions, on page 289](#) for more information on catalog functions.
- Identify the POST request method from the HTTP request and process the parameters and data extracted
- Create a response message to be used in preprocessor or rule or to be sent directly to the HTTP client.

In the event based approach For example, in response to a POST request you might do the following:

- Create a concept instance, using XPath functions to extract data from the POST data in the request event payload
- Create a response event and use `Event.replyEvent()` to send back an empty response using the request event's default destination.

As another example, in response to a GET request you might do the following:

- Identify a concept instance using a property in the request event (created on arrival of the GET request message).
- As needed, identify or generate data to return and create a response event to hold that data.
- Return the data using `Event.replyEvent`.

Task G In the CDD, Configure the Processing Unit

In the Cluster Deployment Descriptor (CDD) Configure the processing unit for deployment as needed (see [Configuring Processing Units \(All OM Types\) on page 513](#) for details). Note the following requirement.

Setting Up Fault Tolerance for the HTTP Channel

Use the Apache HTTP server (“httpd”) and mod_jk to setup fault tolerance to the TIBCO BusinessEvents HTTP channel and a third party fronting server. These are used in addition to sharing the cache server between the nodes in the cluster.

Procedure

1. Create two new properties files with any suitable name. For example, `Ajp_1.properties` and `Ajp_2.properties`.
2. Add the ajp connector port property in these files. For example,
 In `Ajp_1.properties`:

```
<channel path>.ajp.connector.port=8011
```

 In `Ajp_2.properties`:

```
<channel path>.ajp.connector.port=8012
```
3. Specify each of properties file as a part of the TRA file of individual engines, that is, one properties file per engine.
4. Configure the Apache HTTP server (“httpd”) as a fronting webserver. Add the `worker.properties` file in Apache httpd under the `config` folder. In the `worker.properties` file set the Apache httpd for load balancing. Configure two workers to listen to the above specified ajp ports. Disable one worker (for example, `worker2`) by default, so that no requests are sent to it. Now, configure another worker (that is, `worker1`), so that, in case of failure, it redirects all incoming requests to the disabled worker (that is, `worker2`). This worker, in case of failure, now becomes the active worker.

The `worker.properties` file should contain the following properties:

```
worker.list=balancer
worker.balancer.type=lb
worker.balancer.sticky_session=0
worker.balancer.balance_workers=worker1,worker2

worker.worker1.type=ajp13
worker.worker1.host=localhost
worker.worker1.port=8011
worker.worker1.redirect=worker2

worker.worker2.type=ajp13
worker.worker2.host=localhost
worker.worker2.port=8012
worker.worker2.activation=disabled
```

5. Add Tomcat connectors `mod_jk.so` (you can download `mod_jk` from, if you don't have one) under the `modules` folder.

You can download `mod_jk` from the following location:

<http://tomcat.apache.org/download-connectors.cgi>

6. Configure Apache `httpd` to load this module and redirect all requests to the load balancer defined earlier. Add the following snippet to `httpd.conf`, after completion of all the basic module load statements:

```
LoadModule jk_module modules/mod_jk.so

<IfModule jk_module>
JkWorkersFile conf/worker.properties
JkLogFile logs/mod_jk.log
JkLogStampFormat "[%b %d %Y - %H:%M:%S] "
JkRequestLogFormat "%w %V %T"
JkLogLevel info

JkOptions +ForwardKeySize +ForwardURISCompat -ForwardDirectories

<Location //WEB-INF/>
deny from all
</Location>

JkMount /* balancer
</IfModule>
```

7. Start Apache server. Start both the instances of be-engines running the HTTP based project.
8. Open the browser/client and enter the url `http://localhost/<context-path>`, without any port.

The request is now routed to the active worker (`worker1`) only, since the other worker (`worker2`) is disabled. You can verify through logs that logs are being created only for one be-engine.

9. Now shutdown the active be-engine associated with the active worker (`worker1`). All subsequent requests from the client now seamlessly routed to other worker (`worker2`).

To verify check the logs in the corresponding be-engine instance.

HTTP Channel Advanced Configuration Settings

The following settings configure the internal Tomcat HTTP server used by the channel. They are set in the Advanced tab of the HTTP channel resource editor.

Table 20 HTTP Channel Advanced Configuration Settings (Sheet 1 of 7)

Property	Notes
Server Type	The only server type in this release is TOMCAT
Debug Request Info	Flag that identifies if the information is displayed for request debugging.
Connection Timeout (msec)	<p>The number of milliseconds the HTTP server waits after accepting a connection, for the request URI line (that is, the first part of the request message) to be presented.</p> <p>The value 0 means no timeout.</p> <p>The default is 60000.</p>
Accept Count	<p>The maximum queue length for incoming connection requests when all request processing threads are in use. Any requests received when the queue is full are refused.</p> <p>The value -1 means that a default of 10 is set.</p> <p>The default is -1.</p>
Socket Output Buffer Size	<p>The size in bytes of the buffer to be provided for socket output buffering.</p> <p>The value -1 means that the use of a buffer is disabled.</p> <p>The default is 9000.</p>

Table 20 HTTP Channel Advanced Configuration Settings (Sheet 2 of 7)

Property	Notes
Max Processors	<p>The maximum number of simultaneous requests that can be handled by the HTTP server.</p> <p>Set to one of these:</p> <ul style="list-style-type: none"> • A value of 10 or greater. Values of less than 10 are treated as 10. • The value -1 which means that the value of 200 is used. <p>The default is -1.</p>
Connection Linger	<p>The number of milliseconds during which the sockets used by the HTTP server linger (that is, not complete immediately) when they are closed. Use of socket linger allows time for a graceful shutdown sequence to complete.</p> <p>To disable use of socket linger, set to -1.</p> <p>The default is -1.</p>
Enable DNS Lookups	<p>Set to <code>true</code> if you want the calls to <code>request.getRemoteHost()</code> to perform DNS lookups and return the actual host name of the remote client. Set to <code>false</code> to skip the DNS lookup and return the IP address as a string instead (thereby improving performance).</p> <p>By default, DNS lookups are disabled.</p>
TCP No Delay	<p>If the checkbox is checked, the <code>TCP_NO_DELAY</code> option is not set on the server socket.</p> <p>If the checkbox is unchecked, the <code>TCP_NO_DELAY</code> option is set on the server socket. Using <code>TCP_NO_DELAY</code> improves performance under most circumstances.</p> <p>The default is checked.</p>

Table 20 HTTP Channel Advanced Configuration Settings (Sheet 3 of 7)

Property	Notes
Compression	<p>Compression can be used to save server bandwidth. Uses HTTP/1.1 GZIP compression.</p> <p>The compression option set here is used together with the MIME types shown in the Compressible Mime Types setting. Allowable values are:</p> <p>off Disables compression. Values in the Compressible Mime Types setting are ignored.</p> <p>on Enables compression</p> <p>force Forces compression.</p> <p>An integer Specifies a threshold amount of data above which output is compressed. The unit is bytes. For example, if set to 2048 then any file above 2MB will be compressed. If content length is unknown, the output is always compressed. Compression applies to MIME types shown in the Compressible Mime Types setting that are over the threshold (or unknown).</p> <p>The default is off.</p>
Use Body Encoding for URI	<p>If checked, the encoding specified in the <code>contentType</code> HTTP header is used to decode the request URI.</p> <p>If unchecked, the value in URI Encoding field (or its default value) is used.</p> <p>If the URI Encoding setting is specified, then the Use Body Encoding for URI checkbox is ignored.</p> <p>See Sending Non-ASCII Content to Event Properties on page 100 for more details.</p> <p>For an example showing <code>contentType</code>, see Working with Outgoing SOAP Messages (Event Payloads) on page 133.</p> <p>The default is unchecked.</p>

Table 20 HTTP Channel Advanced Configuration Settings (Sheet 4 of 7)

Property	Notes
URI Encoding	<p data-bbox="342 295 1285 355">Specifies the character encoding used to decode the request URI. If not specified, UTF-8 is used.</p> <p data-bbox="342 373 1071 407">If specified, the Use Body Encoding for URI setting is ignored.</p> <p data-bbox="342 425 592 460">The default is UTF-8.</p> <p data-bbox="342 477 1299 512">See Sending Non-ASCII Content to Event Properties on page 100 for more details.</p>
Max KeepAlive Requests	<p data-bbox="342 590 1313 651">The maximum number of HTTP requests that can be pipelined until the connection is closed by the server.</p> <p data-bbox="342 668 1313 737">A value of 1 disables HTTP/1.0 <code>keep-alive</code>, as well as HTTP/1.1 <code>keep-alive</code> and pipelining.</p> <p data-bbox="342 755 1220 815">A value of -1 allows an unlimited amount of pipelined or keep-alive HTTP requests.</p> <p data-bbox="342 833 614 868">The default value is -1.</p>
Max HTTP Header Size	<p data-bbox="342 946 1285 980">The maximum size of the request and response HTTP header, specified in bytes.</p> <p data-bbox="342 998 721 1032">The default is 4096, that is 4 KB.</p>
Max HTTP Post Size	<p data-bbox="342 1119 1299 1215">The maximum size of the POST data, specified in bytes, which is handled by the container FORM URL parameter parsing. You can disable the limit by setting this to less than or equal to 0.</p> <p data-bbox="342 1232 771 1267">The default is 2097152, that is, 2 MB.</p>

Table 20 HTTP Channel Advanced Configuration Settings (Sheet 5 of 7)

Property	Notes
Max HTTP Save Post Size	<p>The maximum size of the POST data, specified in bytes, which is saved or buffered by the container during FORM or CLIENT-CERT authentication. For both types of authentication, the POST data is saved or buffered before the user is authenticated.</p> <p>For CLIENT-CERT authentication, the POST data is buffered for the duration of the SSL handshake and the buffer emptied when the request is processed.</p> <p>For FORM authentication the POST is saved while the user is redirected to the login form and is retained until the user successfully authenticates or the session associated with the authentication request expires.</p> <p>A value of -1 means no limit.</p> <p>A value of 0 disables the saving of POST data during authentication.</p> <p>The default is 4096, that is 4 KB.</p>
Protocol	<p>Specifies the HTTP protocol to use. Options are as follows:</p> <p>http HTTP protocol.</p> <p>https For secure communications (SSL).</p> <p>memory Memory protocol.</p> <p>ajp Apache JServ Protocol, a binary protocol.</p> <p>Default is <code>http</code>.</p>
Max Spare Threads	<p>The maximum number of unused request processing threads that are allowed to exist until the thread pool starts stopping the unnecessary threads.</p> <p>The default is 50.</p>
Min Spare Threads	<p>The number of request processing threads that are created when this Connector is first started. The connector also makes sure that it has the specified number of idle processing threads available. Set this to a value smaller than that set for Max Spare Threads.</p> <p>The default is 4.</p>

Table 20 HTTP Channel Advanced Configuration Settings (Sheet 6 of 7)

Property	Notes
Compressible Mime Types	<p>This is a comma-separated list of MIME types for which HTTP compression may be used.</p> <p>This setting works with the Compression setting. See notes for that setting for more details.</p> <p>The default value is <code>text/html, text/xml, text/plain</code>.</p>
Restricted User Agents	<p>Used to limit support to specific browsers and versions of specific browsers. Comma-separated list containing one or more browser names, such as Mozilla, Internet Explorer. Prepend version with a slash, for example, <code>Mozilla/4.0</code>.</p>
Document Root	<p>The absolute path where static HTML files are stored. The HTTP server retrieves pages from this location. This setting is used by internal TIBCO BusinessEvents applications and may have limited application for other purposes.</p> <p>Tip To enable clients of a TIBCO BusinessEvents HTTP server to view the imported concrete WSDL, provide the URL to the document root folder.</p> <p>No default value.</p>
Document Page	<p>The name of the default static HTML file stored in the document root. This setting is used by internal TIBCO BusinessEvents applications and may have limited application for other purposes.</p>
SSL Server: Key Manager Algorithm	<p>The key manager algorithm for the SSL Service provider.</p> <p>The default is <code>SunX509</code>.</p>
SSL Server: Trust Manager Algorithm	<p>The trust manager algorithm for the SSL Service provider.</p> <p>Ensure that a provider is available for any algorithm other than <code>PKIX</code>.</p> <p>The default is <code>PKIX</code>.</p>

Table 20 HTTP Channel Advanced Configuration Settings (Sheet 7 of 7)

Property	Notes
SSL Server: Protocols	<p>The SSL protocols that can be enabled on the server. Add each protocol using a comma-separated list.</p> <p>The default protocols are the ones supported by the SSL Provider. (SSLv3 and TLSv1 are the most widely supported.)</p> <p>Leave blank to use the default protocols for your SSL provider.</p>
SSL Server: Ciphers	<p>The cipher suites (SSL protocols) used by the server. Add each suite name on a separate line.</p> <p>The following cipher suites are supported:</p> <pre data-bbox="299 708 828 1085"> SSL_RSA_WITH_RC4_128_MD5 SSL_RSA_WITH_RC4_128_SHA TLS_RSA_WITH_AES_128_CBC_SHA TLS_DHE_RSA_WITH_AES_128_CBC_SHA TLS_DHE_DSS_WITH_AES_128_CBC_SHA SSL_RSA_WITH_3DES_EDE_CBC_SHA SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA SSL_RSA_WITH_DES_CBC_SHA SSL_DHE_RSA_WITH_DES_CBC_SHA SSL_DHE_DSS_WITH_DES_CBC_SHA SSL_RSA_EXPORT_WITH_RC4_40_MD5 SSL_RSA_EXPORT_WITH_DES40_CBC_SHA SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA </pre> <p>Leave blank to use the default cipher suites for your SSL provider.</p>

Defining Event Properties for Standard HTTP Header Properties

You need to define corresponding event properties for HTTP header names so that the header names are mapped to the event properties at run time.

Standard HTTP header properties can have a dash (-) in their names. While defining the corresponding event properties for such header properties, use an underscore character (_) instead of a dash(-). A dash is *not* allowed in the event property names.

Following is the list of the HTTP header properties. In this a dash is replaced by an underscore while converting from HTTP headers to Event Properties, and vice-versa.

accept-charset	max-forwards	proxy-authenticate
accept-encoding	proxy-authorization	retry-after
accept-language	user-agent	set-cookie
accept-ranges	content-encoding	transfer-encoding
cache-control	content-disposition	proxy-authenticate
content-type	content-language	retry-after
if-match	content-location	set-cookie
if-modified-since	content-md5	transfer-encoding
if-none-match	content-range	www-authenticate
if-unmodified-since	last-modified	

Using HTTP Functions to Configure HTTP Request Messages

This section explains how HTTP functions are used to send secure and non-secure HTTP requests to other servers, and work with responses received.

HTTP functions are located in the HTTP section of the Standard function catalog.

See [HTTP Channel Advanced Configuration Settings on page 108](#) for details about setting up TIBCO BusinessEvents as a secure server.



For configuring SSL over HTTP when the Identity Resource is of the type Certificate/KeyURL, you need to setup Native Tomcat libraries. See the section Adding Installation or Build Script for OpenSSL for Tomcat in the *TIBCO BusinessEvents Installation* guide.

Generating a Self-Signed SSL Certificate (Keystore)

HTTPS requires use of an SSL certificate. It creates a keystore file and password for the specified type of SSL certificate (keystore).



It is not recommended for production to use self-signed certificates.

If you use trusted certificates, see [Loading Trusted Certificates on page 117](#).

Signature Object createKeystore(String ksFilePath, String ksType, String ksPassword)

Description Creates and returns a keystore object, using the given parameters.

Parameters

Name	Type	Description
ksFilePath	String	The absolute path of the keystore file.
ksType	String	The type of keystore. JKS and PKCS12 are supported.
ksPassword	String	Obfuscated password for the keystore.

Returns

Type	Description
Object	The keystore object

Getting POST Data

Signature `Object getPostData(SimpleEvent event)`

Description Returns the POST data sent in an HTTP POST request (event).

Parameters	Name	Type	Description
	event	SimpleEvent	Request event sent on the HTTP destination. Must not be null.

Returns	Type	Description
	Object	The POST data

Loading Trusted Certificates

As explained in [Task A, Add an HTTP Connection, on page 102](#), ensure that certificates from trusted certificate authorities are stored in the project, using an Identity resource in the TIBCO Shared Resources folder. Alternatively use the `BE_GLOBAL_TRUSTED_CA_STORE` global variable to store a location of the certificates outside of the project.

Signature `Object loadTrustedCertificates(String trustedCertsFolder, String passwordToSet)`

Description HTTPS requires use of an SSL certificate. This function loads a trusted certificate (that is, creates and returns a keystore object) from the trusted certificates folder.

Parameters	Name	Type	Description
	trustedCertsFolder	String	The project path to the folder containing the certificates. If trusted certificates are stored outside the project, use the following construct in the function: <code>System.getGlobalVariableAsString("BE_GLOBAL_TRUSTED_CA_STORE")</code>
	passwordToSet	String	Obfuscated password for the keystore.

Returns

Type	Description
Object	The certificate's keystore object

Sending an Event

To send an event as a request encapsulating request headers as properties.

Signature

```
Event sendRequest (String url, SimpleEvent requestEvent, String responseEventURI, long timeoutMillis)
```

Description

This function sends an event as a response to the request.

Parameters

Name	Type	Description
url	String	The URL for the endpoint that will receive this request.
requestEvent	SimpleEvent	The event to serialize and send to the server.
responseEventURI	String	The fully-qualified path of an event. This event is created when the response is received.
timeoutMillis	Long	The timeout interval for the operation. If the value is -1, the server waits forever.

Returns

Type	Description
Event	An event as a response to the request

Sending an Asynchronous Request (Not Secure)

To send requests to an HTTP server asynchronously, use this function.

Signature

```
String sendAsynchronousRequest(String url, SimpleEvent requestEvent, String correlationId, String successCallbackRuleFunctionURL, String errorCallbackRuleFunctionURL, String methodType)
```

Description

Sends a request to the server specified by the url parameter. When TIBCO BusinessEvents receives a response, the callback function is called.

Returns a correlation ID, which is either passed as input, or is generated from the server if the parameter is null. This ID enables you to correlate a request with its response.

Parameters

Name	Type	Description
url	String	The URL for the server that will receive this request.
requestEvent	SimpleEvent	The event to serialize and send to the server.
correlationID	String	An optional ID to correlate the request and the response. If not specified, the ID is generated by the server.
successCallbackRuleFunctionURL	String	The fully-qualified path of a rule function to be invoked for success case. This rule function is called when a successful response is received. The response event would contain the correlation ID. The rule function must have correlation ID, RequestEvent, and ResponseEvent as parameters.
errorCallbackRuleFunctionURL	String	The fully-qualified path of a rule function to be called in case of an error. This rule function is called when the response received indicates an error. The response event would contain the correlation ID. The parameter can be null. The rule function must have correlation ID, RequestEvent, and ResponseEvent as parameters.
methodType	String	The HTTP method type. Valid values are: GET or POST.

Returns

Type	Description
String	A correlation ID

Sending a Secure Asynchronous Request

To send asynchronous requests using SSL, use this function.

Signature `String sendSecureAsynchronousRequest(String url, SimpleEvent requestEvent, String correlationID, String callbackRuleFunctionURL, String methodType, Object clientIdKeystore, String clientIdPassword, Object trustedCertsKeystore, String trustedCertsPassword boolean verifyHostName)`

Description This function is the same as the `sendAsynchronousRequest()` function, with the addition of the SSL-related parameters, shown below.

Parameters All of the parameters for `sendAsynchronousRequest()` plus the following:

Name	Type	Description
<code>clientIdKeystore</code>	Object	The keystore object for client identity.
<code>clientIdPassword</code>	String	Password for the client ID keystore.
<code>trustedCertsKeystore</code>	Object	Keystore Object for trusted certificates.
<code>trustedCertsPassword</code>	String	Password for the trusted certificates keystore.
<code>verifyHostName</code>	Boolean	Flag for checking if a host name matches the names stored in the server's certificates.

Returns

Type	Description
String	A correlation ID

Sending a Secure Synchronous Request

To send synchronous requests using SSL, use this function.

Signature `Event sendSecureRequest(String url, SimpleEvent requestEvent, String responseEventURI, Object clientIdKeystore, String clientIdPassword, Object trustedCertsKeystore, String trustedCertsPassword, boolean verifyHostName, long timeoutMillis)`

Description This function is the same as the `sendSecureAsynchronousRequest()` function, except that the requests are synchronous. The parameters are shown below.

Parameters

Name	Type	Description
url	String	The URL for the server that will receive this request.
requestEvent	SimpleEvent	The event to serialize and send to the server.
responseEventURI	String	The fully-qualified path of an event. This event is created when the response is received.
clientIdKeystore	Object	The keystore object for client identity.
clientIdPassword	String	Password for the client ID keystore.
trustedCertsKeystore	Object	Keystore Object for trusted certificates.
trustedCertsPassword	String	Password for the trusted certificates keystore.
verifyHostName	Boolean	Flag for checking if a host name matches the names stored in the server's certificates.
timeoutMillis	Long	The timeout interval for the operation. If the value is -1, the server waits forever.

Returns

Type	Description
Event	An event as a response to the request

Configuring TIBCO BusinessEvents as a SOAP Server and Client

Using an HTTP channel and a destination configured to use a SOAP Serializer, TIBCO BusinessEvents can act as a web services platform, sending and receiving SOAP requests. You can configure the project manually or you can use a WSDL import utility, which simplifies configuration.

This release supports SOAP version 1.1. Some understanding of SOAP protocol is required in order to work with this feature. See <http://www.w3.org/TR/soap/> for details.



How Clients View the WSDL When TIBCO BusinessEvents acts as a web service server, clients of the service can view the exported concrete WSDL for the web service using the URL to the document root folder (see [HTTP Channel Advanced Configuration Settings on page 108](#) for details about the document root setting). You must place the WSDL file in that location.

Overview of SOAP Related Resources

The following project resources must be configured so that TIBCO BusinessEvents can send and receive SOAP messages.

SOAP Destinations

When a correctly configured HTTP destination receives a SOAP request, the SOAP serializer deserializes the SOAP message to its corresponding event. That event has to be inherited from a `SOAPEvent`. The event payload contains the SOAP envelope.

SOAP Events

To create a SOAP event, create a `SimpleEvent` that inherits a `SOAPEvent`. It makes event configuration easier. Its payload has a `message` root element having an `Envelope` child element. The root element contains `Header` and `Body` elements, and the `Body` element has a `Fault` element. You can further configure these elements using the payload editor.

SOAP messages (and events) can have attachments.



For outbound SOAP events, ensure that the default destination is specified. If it is not specified, an incorrect SOAP message, with “Message” as its root instead of “Element”, is sent out.

Also, ensure that the default destination for the SOAP event uses `com.tibco.cep.driver.http.serializer.SOAPMessageSerializer`.

SOAP Encoding

Only the SOAP document/literal form of encoding is supported.

Rules and Rule Functions

A rule in the project accesses the SOAP event payload as needed using the mapper or SOAP catalog functions. The rule puts the SOAP response into the payload of a `SOAPEvent` that is sent to the client. The SOAP serializer sends the SOAP response to the client.

SOAP Catalog Functions

SOAP catalog functions enable you to access and process the contents of the following elements in the event payload of an incoming `SOAPEvent`, and add elements to the outbound `SOAPEvent`:

- Header
- Body
- Fault
- Attachment

See [Parsing and Building SOAP Messages on page 130](#) for more details.

Mapping of SOAP Request URI to Destination

For SOAP requests, the header property `SOAPAction` represents the destination name, and the `requestURI` represents the channel URI. TIBCO BusinessEvents combines the `SOAPAction` and `requestURI` to create a destination URI. It maps the request with a destination having same URI as the newly created URI.

For example, if the `requestURI` of a request is `/QueryBooks` and the value of `SOAPAction` is `QueryBooksByAuthor`, then TIBCO BusinessEvents engine maps the request with a destination having `/QueryBooks/QueryBooksByAuthor` URI if it exists.

Manually Creating Resources to Work with SOAP Services

This section explains how you can manually create and configure project resources to work with SOAP services

The section [Creating Resources Using the WSDL Import Utility on page 125](#) explains how to use the WSDL import feature to create the resources.

Also see [Understanding the WSDL to Project Resource Mapping on page 135](#) for details about how the WSDL elements correspond to project resources.

Task A Configure the HTTP Channel and SOAP Destination

Configure an HTTP channel, an HTTP connection resource for the channel. See [Working with HTTP Requests on page 99](#) for details. Connection configuration is the same for SOAP and HTTP channels, except for the serializer. For SOAP destinations, use the following serializer:

```
com.tibco.cep.driver.http.serializer.SOAPMessageSerializer
```

After you complete [Task B](#), set that `SOAPEvent` as the default event for the destination.

Task B Add a SOAPEvent (and Other Ontology as Needed)

Configure the SOAPEvents that will receive the SOAP requests and send out SOAP responses. Also configure any other ontology as needed. `SOAPEvent` is an event type provided with TIBCO BusinessEvents. However it is created in a two-step manner:

1. Add a simple event in the usual way. See [Adding a Simple Event on page 165](#).
2. In the Inherits From field, select **SOAPEvent**.
3. Set the default destination to the destination you created in [Task A](#).

The payload of a `SOAPEvent` is automatically configured with the structure of a SOAP message. It has Header and Body elements, and within the Body element, a Fault element. You can further configure the Header and Body elements using the payload editor.

Task C Configure Rules and Rule Functions using SOAP Functions

In general the procedure of serving requests and sending requests to other servers is the same for SOAP services as for other HTTP interactions. See [Working with HTTP Requests on page 99](#). One rule function is used as the event preprocessor.

In addition, you must configure rules and rule functions to access information from the SOAP messages in the inbound events, and to populate outbound events with SOAP message details. Set two SOAPEvents you configured in [Task B](#) as input and return arguments of the preprocessor. See [Parsing and Building SOAP Messages on page 130](#) for details.

Task D Configure an Event Preprocessor

Open the project CDD for editing, and configure an event preprocessor for the destination you configured in [Task A](#). Use the rule function you configured in [Task C](#).

Creating Resources Using the WSDL Import Utility

WSDL (Web Services Description Language) is an XML format for describing web services as a set of endpoints operating on messages containing either document-oriented or procedure-oriented information.



- Only document style WSDLs with literal encoding are currently supported.
- Only In-Only and In-Out Message Exchange Patterns (MEPs) are currently supported.

TIBCO BusinessEvents can parse an imported WSDL file and create the required project artifacts based on its contents: the channels and destinations (concrete WSDLs only), events, rules, and rule functions.

Both abstract and concrete WSDL files can be used. When you import an abstract WSDL, channel and transport information is not available and you must create those resources manually. Concrete WSDLs contain information that is used to create these resources automatically.

You then implement the newly created rules and rule functions to provide the expected behavior for the web service described in the WSDL file.

See [Understanding the WSDL to Project Resource Mapping on page 135](#) to understand which sections of the WSDL are used to create each resource.

To Import a WSDL File

First you must import the WSDL file into the project. Then you can use the WSDL import utility.

1. Import the WSDL file into the project, as follows.
 - a. In Studio Explorer, right-click the project folder (or the project root) where you want to store the WSDL file and select **Import**. You see the Import wizard.
 - b. At the Select dialog, select **General > From File System**. Click **Next**.
 - c. At the File system dialog, click **Browse**, select the directory that contains the WSDL file you want to use, and click **OK**.
 - d. The File system dialog displays the directory in the left panel and its files in the right panel. Check the checkbox to select the file you want to use. Click **Finish**.

The WSDL file is imported into the project.

2. Create project artifacts from the WSDL file as follows.
 - a. In Studio Explorer, right-click anywhere in the project and select **Import**. You see the Import wizard.
 - b. At the Select dialog, select **TIBCO BusinessEvents > WSDL**. Click **Next**.
 - c. At the WSDL Import dialog, click **Browse** and select the WSDL file you imported in [step 1](#), then click **OK**.
 - d. The WSDL Import dialog displays again, showing the selected WSDL. Click **Finish**.

If the WSDL file you select is an abstract WSDL, you see a warning letting you know that a channel will not be created. See the section introduction for more details about abstract WSDL files.

The project artifacts are created within a folder named using part of the WSDL file name.

3. Save the project.

Re-importing a WSDL File

When you import a WSDL file, TIBCO BusinessEvents creates the project artifacts. If any changes are made to the WSDL file after the import, you must perform the following steps to use the updated WSDL:



Re-importing a WSDL file fails if the project artifacts are already present.

1. Navigate to the project directory and rename the WSDL file.

2. Copy the updated WSDL file to the same directory. Ensure that the name of the WSDL file is the same.
3. In Studio Explorer, select the WSDL file and then select **Edit > Delete** from the top menus to delete the WSDL file.
4. Refresh the project (keyboard shortcut **F5**) and then select **Project > Clean** from the top menu.
5. If you have configured mappings in the project, check for errors caused by broken links in the mapper.
6. Open the mapper and click on the **Mapper Check and Repair** button and check only the mappings in red.
7. Click **Ok** to fix the mappings.
8. If there are new or changed elements, you may have to manually map the new or changed elements.

Exporting Resources as a WSDL File

The WSDL export utility allows you to export rules and rule functions that have a SOAPEvent as a parameter, as a WSDL file.



TIBCO BusinessEvents supports WSDL 1.1 specification. See the following page for more details:

<http://www.w3.org/TR/wsdl>

WSDL Filenames must conform to the `NCName` datatype. See the following page for more details:

<http://www.w3.org/TR/REC-xml-names/#NT-NCName>

Some Japanese characters, such as half-width Katakana, have issues when they are used in XML names. See the following document for more details:

<http://www.w3.org/Submission/japanese-xml/>

[Table 21](#) shows how related project resources are exported as a WSDL file.

Table 21 Exporting Project Artifacts as WSDL

Project Folders and Resources	WSDL Artifact after Exporting
<p>RuleFunction</p> <p>Use the rule function specified as the input event's default destination as the SOAPEvent's preprocessor.</p> <p>That rule function must have the "SOAPEvent for input" as the input and the "SOAPEvent for output" as the return type.</p>	<p>Forms the <wsdl:operation> operation.</p>
<p>SOAPEvent</p> <ul style="list-style-type: none"> • SOAPEvent for the WSDL input argument <p>OR</p> <ul style="list-style-type: none"> • SOAPEvent for the WSDL output argument <p>Use the event you configured in Task B.</p>	<p>Creates a message. The contents of Envelope > Header > Body become message parts.</p>
<p>HTTP Channel</p> <p>Channel with HTTP as the Driver, Resource as the Method of Configuration, with a pointer to the HTTP Connection.</p>	<p>The URI of the channel forms the <soap:address> URI.</p>
<p>HTTP Connection for HTTP Channel</p> <p>Host and port are used.</p>	<p>The name of the connection forms the <wsdl:port>. Host and port of the connection form the <soap:address> host and port.</p>
<p>Destination for HTTP Channel</p> <p>With SOAPMessageSerializer as the default serializer.</p> <p>The destination used is the default destination of the SOAPEvent.</p>	<p>Forms the soapAction attribute of the operation.</p> <p>The destination name becomes soapAction specified for that particular operation.</p>

Table 21 Exporting Project Artifacts as WSDL (Cont'd)

Project Folders and Resources	WSDL Artifact after Exporting
<p>CDD</p> <p>In the Agent Classes tab, Input Destination Collections list, the Preprocessor specified for the input SOAPEvent's default destination.</p> <p>Place the input destination directly under the agent class input destination collections, and not under Collections > Input destination.</p>	<p>Associates <code>soapAction</code> to the operation in the HTTP binding.</p>

To Export a Rule or Rule Function as a WSDL

1. Right-click the project, and click **Export...**
2. In the Export dialog, select **WSDL** under TIBCO BusinessEvents, and click **Next**.
3. Select the WSDL location by specifying the path, and type a valid NCName for the WSDL file.
4. Select the CDD file of the project.
5. Click **Finish**.
6. Save the project.

Parsing and Building SOAP Messages

This section explains how to use the SOAP functions to parse information from incoming requests, and to construct outgoing messages (responses and requests) in your rules and rule functions.

This manual does not explain how to work with the SOAP protocol. For example, you should understand how the `Actor` and `MustUnderstand` attributes in SOAP headers are used to process the message as it passes from its originator, through intermediary applications, to its ultimate destination.

Working with Incoming SOAP Messages (Event Payloads)

An incoming SOAP message can be a request, or a response, depending on whether TIBCO BusinessEvents is acting as the server or the client. TIBCO BusinessEvents can also act as an intermediate node along the path of a SOAP message to its ultimate destination. This section explains how to parse (get) information out of the incoming event payload, which contains the SOAP message.

To Parse the SOAP Envelope

```
String getEnvelope(SimpleEvent inSoapEvent)
```

Given a request `SOAPEvent`, this function returns the SOAP envelope in the request event payload, for example:

```
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
xmlns:sch="http://www.tibco.com/schemas/SoapOverHttp/Schema/Schema.xsd" >
  <soapenv:Header/>
  <soapenv:Body>
    <sch:root>
      <sch:First>1</sch:First>
      <sch:Second>2</sch:Second>
    </sch:root>
  </soapenv:Body>
</soapenv:Envelope>
</message>
```

To Parse (and Optionally Remove) Headers and Header Attributes

```
String[] getHeaders(SimpleEvent inSoapEvent, String actor, Boolean
removeHeaders)
```

If the `actor` parameter has a null value then all the immediate children of the `Header` element are retrieved:

```
getHeaders(inSoapEvent, null, false)
```

Otherwise, the header specified by the actor attribute is retrieved. For example, given this Headers element in a SOAP event payload:

```
<soapenv:Header>
  <t:user xmlns:t="http://schemas.xml.com" soapenv:mustUnderstand="true"
    soapenv:actor="http://localhost:9090/Service">jon</t:user>
  <t:user_surname xmlns:t="http://schemas.xml.com"
    soapenv:mustUnderstand="true"
    soapenv:actor="http://localhost:9090">smith</t:user_surname>
</soapenv:Header>
```

If you specify the following:

```
getHeaders(inSoapEvent, "http://localhost:9090/Service", false)
```

To remove the specified header part or parts, set the final parameter to true. (The SOAP specification states that if a header is processed it should be removed. You would remove a header if TIBCO BusinessEvents is acting as an intermediary node and the request created using the SOAP functions will be sent on to another server.)

Then the first Headers element is returned:

```
<t:user xmlns:t="http://schemas.xml.com" soapenv:mustUnderstand="true"
  soapenv:actor="http://localhost:9090/Service">jon</t:user>
```

You can also retrieve the attributes of a SOAP Header element:

```
String[] getSOAPHeaderAttribute(SimpleEvent inSoapEvent, int index,
String attribute)
```

You can also remove all or selected headers using one of these functions:

```
removeHeaderPart()
removeHeaderParts()
```

To Parse the SOAP Body (SOAPBodyParts)

Two functions are available for getting SOAP body parts.

```
String[] getAllSOAPBodyParts(SimpleEvent inSoapEvent)
```

```
String[] getSOAPBodyParts(SimpleEvent inSoapEvent, String name,
String namespace)
```

The `getAllSOAPBodyParts()` function simply returns all SOAP body parts.

The `getSOAPBodyParts()` function allows you to specify which parts are of interest. Given a body part name and a namespace for a specified `SOAPEvent`, it returns a `String` array of matching SOAP body parts in serialized form. Name and namespace parameters cannot be null.

Example

Given this function:

```
String[] body_part= getSOAPBodyParts
(soapeventin,"root",http://www.tibco.com/schemas/SoapOverHttp/Schema/Schema.xsd)
```

And this `soapeventin` event payload:

```
<?xml version="1.0" encoding="UTF-8"?>
<message>
  <soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
xmlns:sch="http://www.tibco.com/schemas/SoapOverHttp/Schema/Schema.xsd"
xmlns:temp="http://temp/">
  <soapenv:Header/>
  <soapenv:Body>
    <sch:root>
      <sch:First>1</sch:First>
      <sch:Second>2</sch:Second>
    </sch:root>
    <sch:parent2>
      <sch:child1>3</sch:child1>
      <sch:child2>4</sch:child2>
    </sch:parent2>
  </soapenv:Body> </soapenv:Envelope>
```

You would get this as the SOAP body part:

```
body_part[0]=
<?xml version="1.0" encoding="UTF-8"?>
<ns0:root xmlns:ns0="http://www.tibco.com/schemas/SoapOverHttp/Schema/Schema.xsd">
  <ns0:First>1</ns0:First>
  <ns0:Second>2</ns0:Second>
</ns0:root>
```

To Parse Attachments

The following functions enable you to work with SOAP attachments in the request message:

```
getNumberOfAttachments(SimpleEvent inSOAPEvent)
getAttachmentContentID(SimpleEvent inSOAPEvent, int Index)
```

```

getAttachmentContentType(SimpleEvent inSOAPEvent, int Index)
getAttachmentContent(SimpleEvent inSOAPEvent, int Index)
getAttachmentContentByContentID(SimpleEvent inSOAPEvent, string
contentID)

```

The content ID is the attachment identifier. You can select which attachment to work with using its index position. First get the count of the attachments using `getNumberOfAttachments()`. Then using the index, you can get the content ID and content type, as well as the attachment content itself.

The content is returned in byte form, so after you get the content, you must then use other functions to make the content human-readable.

For an example, see

<http://www.w3.org/TR/SOAP-attachments#SOAPReferenceToAttachements>.

To Parse SOAP Fault XML Nodes

The following functions enable you to work with the standard SOAP Fault XML nodes from the payload of a SOAP event:

```

getFault (SimpleEvent, soapEvent)
getFaultActor (SimpleEvent, soapEvent)
getFaultCode (SimpleEvent, soapEvent)
getFaultString (SimpleEvent, soapEvent)

```

Working with Outgoing SOAP Messages (Event Payloads)

You can add header parts, body parts, fault parts and attachments to the outgoing SOAP message (whether it is a response or a request).

The signatures of the relevant functions are as follows:

```

addHeaderPart(SimpleEvent outSOAPEvent, String headerPartXml)
addSOAPBodyPart(SimpleEvent outSOAPEvent, String bodyXML)
addSOAPHeaderAttribute(SimpleEvent outSOAPEvent, int index, String
attribute, String value)
addFaultPart(SimpleEvent outSOAPEvent, String faultCode, String
faultMessage, String faultActor, String faultDetailString)
addAttachment(SimpleEvent outSOAPEvent, String contentID, String
content, String contentType, String contentEncoding)

```

TIBCO BusinessEvents adds each type of fragment to the appropriate part of the event payload — header, body, or fault. The fragments must be well-formed XML. You can also add attachments.

For example, to add a body part containing information for a response you would include all the required details including any namespace information:

```
SOAP.addSOAPBodyPart(outSOAPEvent, "<ns0:BookStore  
xmlns:ns0=\"http://www.abc.com/xsd/books\"><ns0:Book><ns0:Author>J.K.Rowling</ns0:A  
uthor></ns0:Book></ns0:BookStore\"");
```

The specified body part is added to the correct place in the outline structure of the SOAP message, which is provided by the SOAPEvent. The resulting payload would look similar to the following:

```
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/">  
  <SOAP-ENV:Body>  
    <ns0:BookStore xmlns:ns0="HTTP://www.abc.com/xsd/books">  
      <ns0:Book>  
        <ns0:Author>J.K.Rowling</ns0:Author></ns0:BookStore>  
      </ns0:Book>  
    </ns0:BookStore>  
  </SOAP-ENV:Body>  
</SOAP-ENV:Envelope>
```

Understanding the WSDL to Project Resource Mapping

A WSDL file describes a web service. The WSDL Import utility imports a WSDL file and generates TIBCO BusinessEvents project artifacts using elements in the WSDL. TIBCO BusinessEvents can import abstract and concrete WSDL files. The source of the WSDL could be, for example, a ActiveMatrix BusinessWorks SOAPRequestReply activity.



For WSDL import to be successful, all names in the WSDL must conform to TIBCO BusinessEvents folder and entity naming requirements. See [Identifier Naming Requirements on page 323](#).

Example WSDL

The table following this example shows which WSDL elements and attributes are used to create TIBCO BusinessEvents project artifacts. Elements and attributes used in the import are highlighted in bold text. Differences between import from abstract and concrete WSDL files are also highlighted. See [Table 22, Imported WSDL Project Artifacts, on page 136](#) for more details.

```
<?xml version="1.0" encoding="UTF-8"?>
<!--Created by TIBCO WSDL-->
<wSDL:definitions omitted to keep the example short>
  <wSDL:types>
    <xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns="http://www.books.org" elementFormDefault="qualified"
attributeFormDefault="unqualified" targetNamespace="http://www.books.org">
. . . . . Elements omitted
  </xsd:schema>
  </wSDL:types>
. . . . . Elements omitted
```

Note: In an abstract WSDL the following elements are used in the import. However in a concrete WSDL, the `<wSDL:binding>` elements are used instead.

```
<wSDL:portType name="GetBookPortType">
  <wSDL:operation name="GetBook">
    <wSDL:input message="tns:GetBookRequestMessage"/>
    <wSDL:output message="tns:GetBookResponseMessage"/>
  </wSDL:operation>
</wSDL:portType>
```

Note: This is a concrete WSDL example, so the `<wSDL:binding>` elements are used in the import. (In an abstract WSDL, the `<wSDL:portType>` element contents are used instead.)

```
<wSDL:binding name="getBookBinding" type="tns:GetBookPortType">
  <soap:binding style="document"
transport="http://schemas.xmlsoap.org/soap/http"/>
  <wSDL:operation name="GetBook">
```

```

<wsdl:documentation>The operation has no documentation
</wsdl:documentation>
<soap:operation style="document" soapAction="/Service/getBook"/>
<wsdl:input>
  <soap:body use="literal" parts="part1"/>
  <soap:header use="literal"
    message="tns:TransactionRecordMessage" part="user"/>
</wsdl:input>
<wsdl:output>
  <soap:body use="literal" parts="part1"/>
  <soap:header use="literal"
    message="tns:TransactionRecordMessage" part="transactionID"/>
</wsdl:output>
</wsdl:operation>
</wsdl:binding>
<wsdl:service name="getBook">
  <wsdl:port name="getBookHttpPort" binding="tns:getBookBinding">
    <soap:address location="http://ACME:9090/Service/getBook"/>
  </wsdl:port>
</wsdl:service>
</wsdl:definitions>

```

Example Project Folder Structure

Suppose you import the example WSDL above into a Studio project called Library. The imported and generated project artifact names would appear as shown in the Project Folders and Resources column in the table below. Folders that are added by TIBCO BusinessEvents are shown in bold. The example is a concrete WSDL. In the WSDL Source column the source of folder and resource names is given for abstract as well as concrete WSDL sources.

Table 22 Imported WSDL Project Artifacts

Project Folders and Resources	Project Resource Type	WSDL Source
Library	Project root folder	N/A
getBook/	Folder	For concrete WSDLs: <wsdl:service> For abstract WSDLs: <wsdl:service> is not present, so the folder structure starts from the folder created from <wsdl:portType>.

Table 22 Imported WSDL Project Artifacts (Cont'd)

Project Folders and Resources	Project Resource Type	WSDL Source
GetBookPortType/	Folder	For abstract and concrete WSDLs: <wsdl:portType>
getBook/Events/	TIBCO BusinessEvents folder	
GetBookRequestMessage	Event	<wsdl:input> For abstract WSDLs, in <wsdl:portType> section. For concrete WSDLs, in <wsdl:binding> section.
GetBookResponseMessage	Event	<wsdl:output> For abstract WSDLs, in <wsdl:portType> section. For concrete WSDLs, in <wsdl:binding> section.
getBook/RuleFunctions/	TIBCO BusinessEvents folder	
GetBook	Rule function	<wsdl:operation> For abstract WSDLs, in <wsdl:portType> section. For concrete WSDLs, in <wsdl:binding> section.
getBook/Rules/	TIBCO BusinessEvents folder	
GetBookPortType	Folder	Abstract WSDL: <wsdl:portType name> Concrete WSDL: <wsdl:binding type>

Table 22 Imported WSDL Project Artifacts (Cont'd)

Project Folders and Resources	Project Resource Type	WSDL Source
GetBook	Rule	<wsdl:operation> For abstract WSDLs, in <wsdl:portType> section. For concrete WSDLs, in <wsdl:binding> section.
Import from Concrete WSDLs Only If the import is from a concrete WSDL, the HTTP Connection resource, Channel resource and Destination resource are added using details in the <wsdl:service> section of the WSDL. If the import is from an abstract WSDL, you must create these resources manually.		
getbook/Transports/	TIBCO BusinessEvents folder	
getBookHTTPPort	HTTP Connection	<wsdl:port> The host and port come from the <soap:address location>
Service/	TIBCO BusinessEvents folder	
getBook	Channel See Channel Folders on page 138 .	<soap:address location> (from the last part of the location URL)
Service_getBook	Destination See Destination Names on page 139	<soap:operation soapAction>

How Project Artifacts are Named

Channel Folders

Given a concrete WSDL, folders are created for all the elements after the port, up to the last forward slash of the location URL. The text after the last forward slash is the channel name. For example, given the following location URL:

```
http://ACME:9090/Service/Trial/getBook
```

The folder structure would be `/Service/Trial` and the channel name would be `getBook`.

Destination Names

In a concrete WSDL, the `SOAPAction` attribute of a `<soap:operation>` element specifies the URL of a destination. It also becomes the destination name. Forward slashes (/), colons (:), and periods (.) are converted to underscore characters (_) to form the name. For example:

```
http://www.acme.com/TNT/webservices/getByteField
```

Becomes:

```
http___www_acme_com_TNT_webservices_getByteField
```

Rules and Rule Functions

For each operation, the import utility creates a rule and a rule function. The rule has no body. The rule functions have `SoapEventOut` as the return type. Null value is returned by default.

For example, the `GetBook` operation becomes a `GetBook` rule in the `GetBookPortType` folder which is in the `Rules` folder, and also a `GetBook` rule function in the `Rulefunctions` folder.

You implement the rules and rule functions in your project according to the web service you want to implement.

Events

The `<wsdl:input>` element becomes a request event and the `<wsdl:output>` element becomes a response event. Each event type inherits from the `SoapEvent` event type.

Event names come from the `message` attributes. In the example, the request event is `GetBookRequestMessage` and the response event is `GetBookResponseMessage`.

Faults

Faults specified in a WSDL are used in the outbound SOAP event, as the `Fault` element.

Chapter 7 **HAWK Channel**

See [Chapter 4, Channels and Destinations, on page 57](#) for basic information and procedures that apply to all types of channels. This chapter provides additional information about working with the Hawk channel.

Topics

- [Working with the Hawk Channel, page 142](#)
- [Using the Hawk Destination and Event Wizard, page 144](#)
- [Hawk Channel Destination Reference, page 146](#)

Working with the Hawk Channel

The Hawk channel allows TIBCO BusinessEvents to receive events from the Hawk monitor and transform them to events in TIBCO BusinessEvents. Create a channel with driver type Hawk as described in [Adding Channels and Destinations on page 66](#).

Before proceeding to configure the Hawk channel, you must edit the configuration file `studio.tra` at `BE_HOME\studio\eclipse\configuration` and perform the following:

1. Add the following environment variables:

```
tibco.env.HAWK_HOME=<absolute_path_where_TIBCO_Hawk_is_installed>
tibco.env.TRA_HOME=<absolute_path_where_TIBCO_Rendezvous_is_installed>
```

Depending on the transport used by the Hawk Connection shared resource, you must set one of the following:

```
tibco.env.RV_HOME=<absolute_path_where_TIBCO_Rendezvous_is_installed>
tibco.env.EMS_HOME=<absolute_path_where_TIBCO_Enterprise_Message_Service_is_installed>
```

2. Edit the Studio extended classpath and append the following:

```
;%HAWK_HOME%/lib;%RV_HOME%/lib;%TRA_HOME%/lib;%EMS_HOME%/lib
```

3. Restart TIBCO BusinessEvents Studio to pick up the updated values if it is already started.

In order to use the channel to receive events from the TIBCO Hawk monitor, you must configure one or more destinations for the channel. Destinations represent different types of monitors listening to different Hawk events. The following monitor types are available:

- **AlertMonitor** - used to listen on the events for any alerts being posted.
- **RuleBaseListMonitor** - used to listen on the events for changes in the rulebase. For example, if some rulebases are added, updated, or removed.
- **Agent Monitor** - used to listen on the events for any changes in the status of the agents. For example, the status may change from alive to expired.
- **MicroagentListMonitor** - used to listen on the events for any microagents added or removed.
- **ErrorMonitor** - used to listen on the event for errors.
- **WarningMonitor** - used to listen on the event for warnings.
- **Microagent Method Subscription** - used to subscribe to a microagent method at specified time intervals.

Runtime Configuration for the Hawk Channel

Before using the Hawk channel at runtime, edit the configuration file `be-engine.tra` at `BE_HOME\bin` and set the following variables:

1. Add or edit the environment variable `HAWK_HOME`:

```
tibco.env.HAWK_HOME=<absolute_path_where_TIBCO_Hawk_is_installed>
```

2. Depending on the transport used by the Hawk Connection shared resource, set one of the following:

```
tibco.env.RV_HOME=<absolute_path_where_TIBCO_Rendezvous_is_installed>
```

```
tibco.env.EMS_HOME=<absolute_path_where_TIBCO_Enterprise_Message_Service_is_installed>
```

Using the Hawk Destination and Event Wizard

The Hawk Destination and Event Wizard can be invoked from BusinessEvents Studio Explorer, select File > New > Other > TIBCO Channel Wizards > **Hawk Destination and Event**.

You can choose one of the wizard types:

- **Create a Destination and Event** Creates a destination and a simple event for the specified Hawk channel.
- **Create Event for Hawk Catalog Function** Creates a simple event for catalog functions that invoke a specified microagent method from a rule or rulefunction.

Creating a Destination and Event Using the Wizard

1. Start the Hawk Destination and Event Wizard as described in [Using the Hawk Destination and Event Wizard](#) and select the wizard type Create Destination and Event. Click **Next**.
2. On the Select Channel dialog, select the Hawk channel from the list for which you need to create a destination and event. Click **Next**.
3. On the Select Monitor Type dialog, select the monitor type for the Hawk channel. The monitor type determines the type of Hawk events that will be monitored. Click **Next**.

If you selected MicroAgent Method Subscription, proceed to [step 4](#). Else go to [step 7](#).

If you select
MicroAgent
Subscription
Method

4. The wizard detects the available Hawk agents and displays them on the Select Hawk Agent dialog. Select a Hawk agent from the list and click **Next**.
5. The Select Hawk MicroAgent dialog displays. Select a microagent from the list of available Hawk microagents and click **Next**.
6. The Select Hawk MicroAgent Method dialog displays. Select a method from the list of available methods of the microagent. The tooltip displays the arguments for the selected method along with their datatypes.
7. The Set File and Folder name dialog is displayed. Enter the name of the destination to be created, the folder name where the event is to be stored, and the name of the event to be created.
8. Click **Finish**. A message 'Created the resources of Hawk Channel successfully.' indicates that the Hawk destination and the event have been created.

Navigate the project directory in BusinessEvents Studio Explorer to verify that the specified destination and event have been created. The return elements for the selected method are added to the event as properties.

Create Event for Hawk Catalog Function Using Wizard

1. Start the Hawk Destination and Event Wizard as described in [Using the Hawk Destination and Event Wizard](#) and select the wizard type Create Event for Hawk Catalog Function. Click **Next**.
2. On the Select Channel dialog, select the Hawk channel from the list for which you need to create an event. Click **Next**.
3. The wizard detects the available Hawk agents and displays them on the Select Hawk Agent dialog. Select a Hawk agent from the list and click **Next**.
4. The Select Hawk MicroAgent dialog displays. Select a microagent from the list of available Hawk microagents and click **Next**.
5. The Select Hawk MicroAgent Method dialog displays. Select a method from the list of available methods of the microagent.
6. The Set File and Folder name dialog is displayed. Enter the folder name where the event is to be stored and the name of the event to be created.
7. Click **Finish**. A message 'Created the resources of Hawk Channel successfully.' indicates that the event has been created.

Navigate the project directory in BusinessEvents Studio Explorer to verify that the specified event has been created in the specified folder. The return elements for the selected method are added to the event as properties.

Hawk Channel Destination Reference

Field Name	Description
Name	Name of the destination to be created. The default name on the screen is of the format NewDestination_n.
Description	Description of the destination that is to be created.
Default Event	The default event for the destination. You can browse and select an existing event from the project.
Serializer/Deserializer	The default serializer used to deserialize Hawk event to a simple event in TIBCO BusinessEvents and to serialize an event in TIBCO BusinessEvents to a Hawk event. You must configure the serializer (<code>com.tibco.cep.driver.hawk.serializer.HawkSerializer</code>) for every destination created on the Hawk channel.
Monitor Type	Type of monitor listening to different Hawk events. Choose one of: <ul style="list-style-type: none"> • AlertMonitor • RuleBaseListMonitor • AgentMonitor • MicroAgentListMonitor • ErrorMonitor • WarningMonitor • MicroAgentMethodSubscription
Subscription Method URI	URI for the Hawk microagent subscription method. When you use the TIBCO Channel Wizard to create a destination, you can select a subscription method from the list of available methods.
TimeInterval (seconds)	Time interval (specified in seconds) between successive calls to a subscription method.

Field Name	Description
Arguments	<p data-bbox="549 199 1263 291">Optional. The subscription method selected determines if the arguments are required. Some methods do not require any arguments.</p> <p data-bbox="549 314 1228 371">Enter the arguments required by the selected subscription method, if any.</p> <p data-bbox="549 394 1278 482">If you create the destination using the TIBCO Channel wizard, the wizard automatically enters the arguments for the selected subscription method, if any.</p>

Chapter 8 **ActiveSpaces Channel**

See [Chapter 4, Channels and Destinations, on page 57](#) for basic information and procedures that apply to all types of channels. This chapter provides additional information about working with the ActiveSpaces channel.

Topics

- [Basic TIBCO ActiveSpaces Concepts and Terminology on page 149](#)
- [Working with ActiveSpaces Channels, page 155](#)

Basic TIBCO ActiveSpaces Concepts and Terminology

This section provides a brief introduction to the ActiveSpaces concepts and terminology that are useful when working with the ActiveSpaces channel. See *TIBCO ActiveSpaces* documentation for more information.

Metaspace

A *metaspace* is a virtual entity that contains spaces, which are containers that store the data used by applications.

A metaspace is a container for managing system spaces, user spaces, and a group of members that are working together in a cluster. The metaspace is the initial handle to ActiveSpaces. An application or member first joins a metaspace, and through it, gets access to other objects and functionality.

A metaspace is created when the first process connects to it, and disappears when the last process disconnects from it. The metaspace grows or shrinks automatically as members connect to it and disconnect from it.

Initially, a metaspace contains only system spaces. As users create spaces in the metaspace, the definition of those spaces (along with other administrative data) is stored in system spaces.

Multiple metaspaces may exist at the same time, each one containing a completely independent set of spaces. This means, for example, that changes to a space called **clients** in a metaspace named **Dev** have no impact on a space named **clients** in a metaspace named **Prod**. Since no single application can connect to two different metaspaces using the same metaspace name, metaspaces should always use different names.

Space

A space provides shared virtual storage for data. A space is:

- A container for a collection of entries that consist of a tuple and associated metadata.
- A *shared* entity. Many applications can access a space concurrently and each application has the same coherent view of the data contained in the space. The spaces in ActiveSpaces are called tuple spaces, and the items stored in them are called tuples.

A space can proactively notify applications of changes in the data contained in the space as changes happen (push model), and can therefore be used as a coordination mechanism for building distributed systems.

- A *virtual* entity that is distributed and implemented collaboratively by a group of processes located on multiple hosts and communicating over the network.

ActiveSpaces handles changes in this set of processes automatically: processes may join or leave the group at any time without requiring any user intervention. A space automatically scales up as the number of processes in the group increases, and scales down when processes suddenly disappear from the group or network. There is no negative impact on the data contained in the space when processes leave the space.

System spaces are a set of administrative spaces that are created and maintained by ActiveSpaces and are used to describe the attributes of the spaces.

User spaces are spaces that are defined by the user.

Members and Member Roles

Entities that need access to a space join the space as *members*.

Members can join a space as a seeder or as a leech:

- *Seeders* play an active role in maintaining the space by providing CPU and RAM.
- *Leeches* play a passive role. They have access to space data but provide no resources.

Seeders

Seeder applications join the space and indicate their willingness to lend some of the resources of the host where they are deployed to scale the service provided by ActiveSpaces. In effect, the seeding applications have a portion of the Space embedded in their process.

ActiveSpaces distributes the data stored in the space evenly between all of the processes that have joined the space as seeders. ActiveSpaces is an elastic distributed system. Seeders can join and leave the space (effectively scaling it up or down) at any time without the need to restart or reconfigure any other participant in the space. When this happens, the distribution of entries is automatically rebalanced if necessary to maintain even distribution between the seeders.

Leeches

An application can also join a space as a *leech*, without contributing any of its host's resources to the scalability of the space, and without experiencing any limitation in functionality or in its ability to use the space. A leech (as opposed to a seeder) is a peer of the space that is not considered by the distribution algorithm and therefore can join or leave the space without causing redistribution of the data.

ActiveSpacesSerializer

The default serializer used to deserialize ActiveSpaces tuple to a simple event in TIBCO BusinessEvents and to serialize an event in TIBCO BusinessEvents to an ActiveSpaces tuple. You must configure the serializer (`com.tibco.cep.driver.as.serializers.ActiveSpacesSerializer`) for every destination created on the ActiveSpaces channel.

Distribution role

The distribution role ([Seeders](#) or [Leeches](#)) is a level of participation of a space member and does not indicate any limitation on use. Leeches have access to the same set of space operations as seeders. The choice of distribution role must be made on a per space basis: it may be that the best solution is to join some spaces as a seeder and others as a leech.

Consumption mode

ActiveSpaces can proactively notify applications of changes to the tuples stored in a space. You can choose one of the consumption modes: Event Listener or Entry Browser.

When the consumption mode of a destination is *event listener*, the destination behaves like a *subscriber* in a publish-subscribe messaging system. When certain data changes or certain events occur, a callback function is invoked. You can choose the type of events to listen to: Put Event, Take Event, and Expire Event.

When the consumption mode is *entry browser*, the destination can monitor the space for data changes or certain events to occur, and can retrieve the tuple from the space using the *browser type* - Get or Take.



The `DistributionScope` as well as the `TimeScope` properties now have options enabled that can be specified for both `EventListener` and `EventBrowser`.

Tuple

A *tuple* is similar to a row in a database table: it is a container for data. Specifically, it is a sequence of named elements called fields (similar to the columns in a database table) which contain values of a specific type. Each tuple in a space represents a set of related data.

Fields have a name and a type. A tuple can be seen as a kind of map on which fields can be 'put' or 'removed'. A tuple can also be seen as a self-describing message. Tuples are platform independent, and can be serialized and deserialized.

Filters

Filters can be applied to both listeners and browsers, and also be used to evaluate a tuple against a filter. Filters allow your application to further refine the set of tuples it wants to work with using a space browser or event listener.

A filter string can be seen as what would follow the *where* clause in a `select * from Space where...` statement.

Examples

```
field1 < (field2+field3)
state = "CA"
name LIKE ".*John.*" //any name with John
```

Filters can make reference to any of the fields contained in the tuples. Filters do not provide any ordering or sorting of the entries stored in the space.

Operators Supported in Filters

Table 23 shows the operators that are supported in the ActiveSpaces filters:

Table 23 Operators for ActiveSpaces Filters

Operator	Meaning
>, >=	greater than
NOT or ! or <>	not
*	multiply
=	equal
!=	not equal
ABS	absolute value

Table 23 Operators for ActiveSpaces Filters

Operator	Meaning
MOD or %	modulo, as in (MOD (x,y) or x % y)
NOT BETWEEN	not between
BETWEEN	between
	string concatenation
NOT LIKE	not like (regex)
LIKE	like (regex)
<, <=	less than
+	addition
OR	or
-	subtraction
AND	and
IN	range, as in "age in (1,3,5,7,11)"
NULL	does not exist
NOT NULL	exists
IS	only used with NULL or NOT NULL, as in "x IS NULL" or "x IS NOT NULL"
NOR	nor, as in "age NOT 30 NOR 40"
/*	comments
//	comments

Formats for Filter Values

Table 24 shows the formats for values used in filters.

Table 24 *Formats for Filter Values*

Filter Value	Format
octal value	\oXXX
hexadecimal value	\xXXX
exponents (as in 1E10 or 1E-10)	XXXEYY
date time	YYYY-MM-DDTHH:MM:SS
date	YYYY-MM-DD
time	HH:MM:SS:uuuu+/-XXXXGMT
true	TRUE
false	FALSE

Working with ActiveSpaces Channels

Create a channel of type ActiveSpaces as described in [Adding Channels and Destinations on page 66](#). In order to use the channel to monitor a space in ActiveSpaces, you need to add and configure one or more destinations for the channel.

Creating a Destination and Event Using the Wizard

TIBCO Channel Wizards provide an option to create a destination and event for a specific ActiveSpaces channel.



The wizard requires the metaspace and spaces in TIBCO ActiveSpaces to be created and initialized. Otherwise, no spaces are listed on the Select space dialog. You must create and initialize the metaspace and then click **Refresh** to refresh the list of spaces that the channel can connect to.

1. In BusinessEvents Studio Explorer, select the ActiveSpaces channel for which you need to create a destination and event. From the menu, select File > New > Other > TIBCO Channel Wizards > **ActiveSpaces Destination and Event** and click **Next**.
2. The Select Space dialog lists the available spaces for every ActiveSpaces channel in the project.



If no spaces are listed on the Select space dialog, check if the metaspace is created and initialized. Create and initialize the metaspace, and then click **Refresh** to refresh the list of spaces that the channel can connect to.

Select a space and click **Next**.

3. On the New SimpleEvent and Destination dialog, enter the values to configure the simple event and destination for the selected channel and space.

Table 25: Configuring an ActiveSpaces Destination and Event

Field Name	Description
Simple Event	
Name	Name of the simple event to be created. The default name is of the format <code><spaceName>Event</code> .

Table 25: Configuring an ActiveSpaces Destination and Event

Field Name	Description
Destination	
Name	Name of the destination to be created. The default name on the screen is of the format <code><spaceName>Dest</code> .
Default Event	The default event for the destination. This is set to the SimpleEvent specified to be created by the wizard.
Serializer/Deserializer	<p>The serializer used to map tuples in ActiveSpaces to simple events in TIBCO BusinessEvents.</p> <p>The default value is <code>com.tibco.cep.driver.as.serializers.ActiveSpacesSerializer</code>.</p> <p>See ActiveSpacesSerializer on page 151 for more information.</p>
Space Name	<p>Name of the space which the destination connects to. This space is selected in the Select space dialog.</p> <p>See Space on page 149 for more information about spaces in TIBCO ActiveSpaces.</p>
Distribution Role	<p>The level of participation of the space member: seeder or leech.</p> <p>See Distribution role on page 151 for more information.</p>
Filter	<p>String specified to evaluate tuples and refine the set of tuples to work. A filter string can be seen as what would follow the <i>where</i> clause in a <code>select * from Space where...</code> statement.</p> <p>See Filters on page 152 for more information.</p>
Consumption Mode	<p>Specifies the consumption mode for the ActiveSpaces event as one of:</p> <ul style="list-style-type: none"> • Event Listener - listens for specific events to occur, and invokes a callback function from TIBCO ActiveSpaces. • Entry Browser - can listen and retrieve tuples from the space using the Get or Take methods. <p>See Consumption mode on page 151 for more information.</p>

Table 25: Configuring an ActiveSpaces Destination and Event

Field Name	Description
Browser Type	<p>Available only when the consumption mode is Entry Browser.</p> <p>Specifies the browser type used to retrieve tuples from a space. You can choose either <i>Get</i> or <i>Take</i>.</p> <p>The difference between the Get and Take browsers is that Get retrieves a copy of the tuple from the space and Take retrieves the tuple from the space, and there is no trace of the tuple in the space after the Take event.</p>
Put Event	<p>Available only when the consumption mode is Event Listener.</p> <p>When selected, the Event Listener listens for any Put events on the space and invokes a callback function when such an event occurs.</p>
Take Event	<p>Available only when the consumption mode is Event Listener.</p> <p>When selected, the Event Listener listens for any Take events on the space and invokes a callback function when such an event occurs.</p>
Expire Event	<p>Available only when the consumption mode is Event Listener.</p> <p>When selected, the Event Listener listens for any Expire events on the space and invokes a callback function when such an event occurs.</p>

- Click **Finish** to create the destination and the simple event.



The wizard adds the following properties to the simple event:

- **id** of type long
- **consumption_mode** of type string
- **browser_type** of type string
- **event_type** of type string

Configuring the Destination and Default Event Manually

Adding a Destination for the ActiveSpaces Channel

You can also add a destination to the ActiveSpaces channel manually, without using the TIBCO Channel wizards.

The following procedure describes the steps to add and configure a destination for the ActiveSpaces channel.

1. In BusinessEvents Studio Explorer view, double-click the channel name to open the channel editor, if it is not already open.
2. In the Destinations section, click **Add**.

Common Fields

3. Enter a Name and Description for the destination.
4. In the Default Event field, browse to and select the event to be created (by default) from incoming messages received by this destination.

If you have not yet created the event, you can select the default event later. See [Assigning a Default Event to the Destination on page 158](#) for details about the simple event.

5. Select the serializer/deserializer to be used:
`com.tibco.cep.driver.as.serializers.ActiveSpacesSerializer`.
6. Enter the values for the rest of the fields as described in [Table 25; Configuring an ActiveSpaces Destination and Event, on page 155](#).

Assigning a Default Event to the Destination

When creating a simple event manually, ensure that you add the property `id` of type `long`.

To assign a default event to a destination, open the channel editor in BusinessEvents Studio Explorer view. In the Default Event field, browse to and select the simple event created earlier.

Support for ActiveSpaces Security

To configure security for ActiveSpaces:

1. Select “Enable transport security”.

When transport security is not enabled, all the fields underneath are hidden (disabled).

2. Select either Controller or Requestor as a security role.
3. If Controller is selected, browse to select the required Policy file.
4. If Requestor is selected, browse to select the required Token file.

Chapter 9 **Simple Events**

TIBCO BusinessEvents supports three sorts of events: Simple events (usually referred to just as events); time events, which are timers; and advisory events. This chapter explains how to use simple events.

Topics

- [Overview of Simple Events, page 162](#)
- [Working with Events in Rules, page 164](#)
- [Adding a Simple Event, page 165](#)
- [Simple Event Reference, page 167](#)
- [Simple Event Attributes Reference, page 173](#)

Overview of Simple Events

A simple event defines an object that represents an occurrence of something. When the term "event" is used without the qualifier advisory or time, it refers to a simple event.

See also

- [Effect of Cache Only Cache Mode on page 253](#)
- *TIBCO BusinessEvents Getting Started* provides a practical introduction to events, including use of default events and default destinations.
- Chapter 2, Channels and Events in *TIBCO BusinessEvents Architect's Guide* for a more detailed overview and information about the following topics:
 - Message Acknowledgement
 - Default Destinations and Default Events
 - Simple Events — Time to Live and Expiry Actions
- [Working with Concept and Event Properties on page 331](#).
- [Chapter 41, Diagrams, on page 701](#) for information on using event dependency and sequence diagrams, and event model views.

Using Inheritance

Use of inheritance can simplify event configuration. A child event inherits the following:

- All the parent event's properties.
 - The parent event's expiry action (if set). However, an expiry action set in the child event overrides the parent event expiry action.
- 

 - A parent event cannot have a payload.
 - All child events of an event with TTL=0 must also use TTL setting TTL=0

Events that are related to each other directly or indirectly by inheritance cannot have distinct properties that share a common name. Therefore the following restrictions apply:

- If two events are related by inheritance, you cannot create a new property in one with a name that is already used in the other.

- If two unrelated events have properties that share a name, you cannot create an inheritance relationship between the two events.

Working with Events in Rules

This section explains various aspects of using events in rules.

See also [Using Scheduler Functions \(Requires Cache OM\) on page 185](#).

Explicitly Assert Events Created in Rules

At runtime, event instances that are created using rules are not automatically asserted. You must explicitly assert such events, for example using the `Event.assertEvent()` function.

Events that are created from incoming messages, on the other hand, are automatically asserted.

Specifying Default and Non-Default Destinations

TIBCO BusinessEvents includes two functions that allow you to send simple events out to another application: `Event.sendEvent()` and `Event.routeTo()`.

- `Event.sendEvent()` automatically sends the event to its default destination.
- `Event.routeTo` takes a destination as an argument, ignoring the event's default destination.

With `routeTo` you can direct an event to a destination on a different channel from the event's default destination. You can also override the properties of the destination, for example, the subject.

You cannot, however, override the properties of the channel itself, for example, the network field in a Rendezvous channel.

Scheduling Simple Events

You can use two methods to schedule simple events:

- Rule-based time events schedule the assertion of simple events.
- Scheduler functions schedule the sending of simple events to their default channels.

Each method is appropriate in different situations. See [Chapter 10, Time Events and Scheduler Functions, on page 175](#) for details.

Adding a Simple Event

This section provides summary steps for adding a simple event. See [Simple Event Reference on page 167](#) for details on how to complete the values. Also see [Overview of Simple Events on page 162](#).

To Add a Simple Event

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the event and select **New > Simple Event**. You see the New Simple Event Wizard.
2. In the Simple Event Name field, type a name for the event. In the Description field, type a description as desired.



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. Click **Finish**. You see the Simple Event Editor.
4. Complete the Standard tab and Advanced tab sections as explained in [Simple Event Reference on page 167](#).
5. Save the resource.



The ability to add metadata groups and properties is provided for customization purposes and is not documented.

Defining Payloads Using XML Schema Files

You can define the XML schema file within TIBCO BusinessEvents Studio, or import an existing one. Then you can reference the schema when defining the payload.

To access an XML payload in rules you can use the @payload attribute, as well as Event, Mapper, and XPath functions. For simple XML payloads you can use string functions.

To Define an XML Schema File in TIBCO BusinessEvents Studio

An Eclipse resource is available for this task.

1. Right-click the folder where you want to add the resource and select **New > Other > XML > XML Schema** and click **Next**.
2. Enter a name and click **Finish**.
3. In the XML schema editor use the **Design** or **Source** tabs as preferred to define the schema. Refer to Eclipse help for details.

To Import an XML Schema File

An Eclipse resource is available for this task.

1. From the **File** menu select **Import > File System**.
2. At the **Import from File System** wizard, browse to a folder and select the desired schema file within that folder.
3. In the **Into Folder** field, specify the project folder where you want to put the file.
4. Set the **overwrite** and **folder options** as desired, then click **Finish**.

To Use an XML Element in a Schema File in an Event Payload

To use an XML element in a schema file as the root element of the event payload, do the following.

1. In the event's **Advanced** tab, click the **Add (+)** button. You see a root node in the left panel.
2. In the **Content** field, select **XML Element Reference**.
3. In the **Schema** field click the **Browse** (binocular) button.
4. You see the **Select a Resource** dialog. Depending on the XML schema files available in the project, multiple entries may be grouped according to location (on the **By Location** tab) and namespace (on the **By Namespace** tab). Select one XML schema resource from either of these tabs.
5. In the **Element** panel, select the element you want to use as the root element for the payload. (You can also change this in the **Payload** section of the Event's **Advanced** tab.)
6. Click **OK** to return to the **Advanced** tab. The word **root** in the left panel is now replaced with this element name.
7. Save the resource.

Simple Event Reference



Simple Event resources are used to define an object that represents an occurrence, such as sending an invoice, debiting an account, and so on.

You can modify and enrich events before they are asserted into the Rete network. Rule evaluation depends on event values at time of assertion, so they can be changed only before assertion.



If you are working with a project imported from a release earlier than 5.0.0, you may see metadata properties. However, do not use them. Instead use the settings and properties in the Domain Objects section of the CDD file as needed. See [Chapter 28, Domain Objects Configuration](#).

Wizard and Configuration (Standard Tab)

The Wizard and the Configuration section have the following fields.

Field	Global Var?	Description
Name	No	Not shown as a field because it cannot be changed. The name appears in the Wizard, and in the title of the event. The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
Description	No	Short description of the resource.
Inherits From	No	This event inherits from the event you select here. Leave blank if you do not want to use inheritance. As a convenience, you can open the selected event resource by clicking the underlined label.

Field	Global Var?	Description
Time to Live	Yes	<p>Specify a numerical value and a time unit.</p> <p>Time units can be specified from the drop-down list available to the right of the Time to Live field. Choose one of: milliseconds, seconds, minutes, hours, and days. The default time unit is seconds.</p> <p>The numerical value is interpreted as follows:</p> <ul style="list-style-type: none"> • One or higher (>0): the event expires after the specified number of units elapse. • Zero (0): the event expires after the completion of the first RTC cycle. • A negative integer (<0): the event does not expire, and must be explicitly consumed. The value -1 is generally used to indicate an event that does not expire. <p>See Declaration and Expiry Action (Advanced Tab) on page 169.</p>
Default Destination	No	<p>When the destination is not otherwise specified (for example in rules or rule functions), events of this type are sent to the destination you select here. You can send an event to the default destination of its event type using the <code>Event.sendEvent()</code> function. As a convenience, you can open the selected destination resource by clicking the underlined label.</p>
Retry On Exception	No	<p>When a destination's event preprocessor fails due to an exception, the behavior of an event instance of this type is determined by this check box setting:</p> <ul style="list-style-type: none"> • When this checkbox is checked TIBCO BusinessEvents attempts to reprocess the event instance that failed and retries indefinitely with a delay of five seconds between retries. • When the checkbox is unchecked TIBCO BusinessEvents does not attempt to reprocess the event instance that failed. <p>See Event Preprocessors on page 275 for more information.</p>

Properties (Standard Tab)

The Properties section has the following fields. Event properties generally map to incoming or outgoing message properties.

Field	Global Var?	Description
Name	No	<p>The name to appear as the label for the property. Names follow Java variable naming restrictions. Do not use any reserved words. See Identifier Naming Requirements on page 323.</p> <p>Note: In addition to standard naming restrictions, do not begin an event property name with <code>_ns_</code> or <code>_nm_</code>. These have a special use. See Mapping Incoming Messages to Non-default Events on page 61.</p> <p>For events used in JMS channels Names beginning with <code>_jms</code> or <code>jms</code> (case insensitive) are used only for JMS header properties. You can, however, use properties beginning <code>jms_</code> (case insensitive) for event properties.</p> <p>See Using JMS Header Properties in Incoming and Outgoing Messages on page 93 for more details. Table 19, JMS Header Field Names, on page 95 shows the list of JMS header properties. Consult the JMS specification for more details.</p>
Type	No	<p>One of: <code>String</code>, <code>Integer</code>, <code>Long</code>, <code>Double</code>, <code>Boolean</code>, <code>DateTime</code></p> <p>Note: For properties of type <code>Double</code>, all NaN (Not a Number) values are converted to <code>0.00</code>.</p>
Domain	No	<p>Click the search button and select the domain model you want to use for this property. See Chapter 14, Domain Models, on page 207 for details on adding domain models.</p>

Declaration and Expiry Action (Advanced Tab)

If the [Time to Live](#) field is zero or higher, define the action or actions to take when an event expires in the Expiry Action section.

If an event is explicitly consumed in a rule, TIBCO BusinessEvents does not execute the expiry action.

The editor in the Expiry Action section is the same as the Rule function editor. See [Rule Function Resource Reference on page 263](#) for details.

See Simple Events — Time to Live and Expiry Actions in *TIBCO BusinessEvents Architect's Guide* for background details.

Payload (Advanced Tab)

An event can have a payload. The payload often corresponds to a message body. Payloads can be defined using an XML schema. In the left panel you add groups (elements) and parameters (attributes). You can add groups as children of a selected group, or at the same level, to define a hierarchy as desired. In the right panel, you define the type of each element or parameter. The table below describes the payload parameters available for each content type. The content types appear in the drop-down list for the Content field in the Payload section:

Table 26 Simple Event Payload Element Parameters (Sheet 1 of 3)

Content/Parameter	Description
Complex Element	An element that contains other elements. This is like a structure in a programming language. The complex element can contain zero or more elements of other types, including other complex elements.
Name	The name of the element.
Cardinality	Values for Cardinality: <ul style="list-style-type: none"> • Required: The payload must include an instance of this element. • Optional (?); The element is not required. • Repeating (*); The element is a list that has zero or more instances. • At least one (+): The element is a list that has one or more instances.
Element of Type	An element with a specified data type. You can specify a scalar data type (string, integer, and so on), you can reference an XML type, or you can specify the TIBCO <code>ActiveEnterprise</code> Any data type.
Name	The name of the element.
Cardinality	See Cardinality under Complex Element.
Type	The generic data type. For example, decimal or date/time.
Type	The specific data type. For example, float or month. Refer to the <i>TIBCO ActiveMatrix BusinessWorks Palette Reference</i> for a complete list.
XML Element Reference	A reference to an element in a stored XML schema. See TIBCO Designer documentation for more information about XML schemas.
Cardinality	See Cardinality under Complex Element.

Table 26 Simple Event Payload Element Parameters (Sheet 2 of 3)

Content/Parameter	Description
Schema	Stored XML schema that contains the element or type you want to reference.
Element	The element within the stored XML schema that you want to reference.
Attribute of Type	An attribute with a specified data type. You can specify a scalar data type (string, integer, and so on), you can reference an XML type, or specify the TIBCO ActiveEnterprise Any data type.
Name	The name of the element.
Cardinality	See Cardinality under Complex Element.
Type	The generic data type. For example, decimal or date/time.
Type	The specific data type. For example, float or month. Refer to the <i>TIBCO ActiveMatrix BusinessWorks Palette Reference</i> for a complete list.
Sequence	A sequence of elements. Each item in the sequence is a structure of the sub-elements of this element.
Cardinality	See Cardinality under Complex Element.
Choice	A choice of elements. The data type of this element can be one of the sub-elements defined.
Cardinality	See Cardinality under Complex Element.
All	The data type of this element can be all of the data types of the sub-elements defined.
Cardinality	See Cardinality under Complex Element.
XML Group Reference	A reference to an XML group in a stored XML schema. See TIBCO Designer documentation for more information about XML schema.
Cardinality	See Cardinality under Complex Element.
Schema	Stored XML schema that contains the element or type you want to reference.
Model Group	Select the appropriate model group from the pull-down list.

Table 26 Simple Event Payload Element Parameters (Sheet 3 of 3)

Content/Parameter	Description
Any Element	A reference to any XML Element. You can use the Coercions button to supply a reference to the XML Element for this item when it appears in the input or process data.
Cardinality	See Cardinality under Complex Element.
Validation	Select the level of validation to be performed on the XML Element: <ul style="list-style-type: none"> • Strict: Must validate by locating a declaration for the element. • Skip: Do not validate. • Lax: Validate if a declaration for the element is available.

Simple Event Attributes Reference



When using an event in a rule's form editor, type `eventname@` to see the list of its attributes.

You can use the following attributes in rules to return information about a simple event instance.

Attribute	Type	Returns
@id	long	The event's unique internal ID.
@extId	string	<p>The event's unique external ID. Optional.</p> <p>The value of the <code>extId</code> is set at creation time, for example, using the event's ontology function, or the <code>Event.createEvent()</code> function. The value cannot be changed after that.</p> <p>Note: The <code>extId</code> value (if set) must be unique across all objects in the cluster.</p> <p>Tip: You can use the property <code>Agent.AgentClassName.checkDuplicates</code> to check for duplicate <code>extIds</code> across the cluster. See CDD Agent Classes Tab Properties Reference on page 501 for details.</p>
@ttl	long	The event's time to live, where the assertion of the event defines the start of the time to live period. You can specify the value in the SimpleEvent resource TTL field. See Simple Event Reference on page 167 .
@payload	string	The payload as a string value. See Payload (Advanced Tab) on page 170 for more on specifying the payload in a SimpleEvent resource.

TIBCO BusinessEvents supports three sorts of events: Simple events (usually referred to just as events); time events, which are timers; and advisory events. This chapter explains how to use time events.

In addition you can use functions to schedule events. This method can be preferable to using time events in some cases.

Topics

- [Overview of Time Events, page 176](#)
- [Working With Time Events, page 178](#)
- [TimeEvent Resource Reference, page 180](#)
- [TimeEvent Attributes Reference, page 182](#)
- [Rule Based TimeEvent Function Reference, page 184](#)
- [Using Scheduler Functions \(Requires Cache OM\) on page 185](#)

Overview of Time Events

TIBCO BusinessEvents offers two kinds of time events, repeating and rule-based. In addition you can schedule events using functions.



Time events do not go through an event preprocessor. If you are using cache-only cache mode, ensure that any objects are properly loaded. Events scheduled using scheduler functions, however, are sent through channels and would therefore go through event preprocessors in the usual way.

See Also

- [Chapter 9, Simple Events, on page 161](#) for overview information about events in general.
- [Chapter 41, Diagrams, on page 701](#) for information on using diagrams.

Scheduled Time Events



Repeat (Interval) based time events are not supported in multi-engine mode Time events configured to repeat at intervals are not supported in multiple-agent (multi-engine) configurations. Rule-based time events, however, are supported.

You can configure a time event to repeat at a configurable time interval. For example, if you configure a time event to repeat every thirty seconds, then every thirty seconds TIBCO BusinessEvents creates a new time event of that type.

You can configure a repeating time event to create a specified number of events at each interval.

At engine startup The time interval begins during engine startup. See Engine Startup and Shutdown Sequence in *TIBCO BusinessEvents Administration* for specific details.

Rule Based Time Events

A rule based TimeEvent resource has only a name and description. You can then use it in a rule to schedule a simple event to be asserted, using its ontology function, `ScheduleTimeEventName()` in a rule (see [Rule Based TimeEvent Function Reference on page 184](#)). You can schedule the event to be asserted after a period of time, and you can pass information to the event and specify its time to live. You can call the `ScheduleTimeEventName()` function in different places with different time delays.

You can use rule based time events in various ways. For example, you might write rules that check for delays in order fulfillment:

1. A new Order event is asserted, and Rule A (which has Order in its scope) creates a time event T and configures it to be asserted in sixty minutes, and passes the order ID as the closure parameter value. (Rule A also sends the order details to another system.)
2. Sixty minutes after Rule A executes, timer event T is asserted.
3. The assertion of time event T triggers Rule B, which has T in its scope. Rule B checks the order status. If the order is delayed, it sends out an alert.

Working With Time Events

This section explains how to create and configure a time event.

Adding a Time Event

See [TimeEvent Resource Reference on page 180](#) for information on how to complete the values.

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the time event and select **New > Time Event**. You see the New Time Event Wizard.
2. In the Time Event Name field, type a name for the time event. In the Description field, type a description.



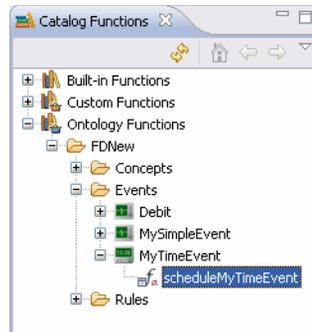
You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. Click **Finish**. You see the Time Event editor.
4. In the Type field, select the type of time event you want to create and configure accordingly:
 - ruleBased: Select Rule Based. Use the event in a rule as explained in [Configuring a Rule Based Time Event in a Rule or Rule Function on page 178](#).
 - repeat: The event occurs at specified intervals. Select Repeat and configure the time event following guidelines in [TimeEvent Resource Reference on page 180](#).
5. Click **Apply** and save the resource.

Configuring a Rule Based Time Event in a Rule or Rule Function

1. Open the rule editor for the rule or rule function where you want to use the rule-based time event.
2. Display the Catalog Functions view if it is not available: Select **Window > Show View > Other > TIBCO BusinessEvents > Catalog Functions**.

3. In the Catalog Functions view, expand Ontology Functions and drill down to the rule-based time event. Expand the rule-based time event and select its ontology function (`scheduleTimeEvent`).



4. Drag the time event's ontology function into the Actions area of the rule (or Body area of the rule function) and configure the parameters, following guidelines in [Rule Based TimeEvent Function Reference on page 184](#).
5. Save the resource.

TimeEvent Resource Reference



TimeEvent resources are used as timers. You can configure rule-based timers, which are scheduled in a rule, and repeating timers, which are scheduled in the TimeEvent resource. Use time events to trigger rules.



If you are working with a project imported from a release earlier than 5.0.0, you may see metadata properties. However, do not use them. Instead use the settings and properties in the Domain Objects section of the CDD file as needed. See [Chapter 28, Domain Objects Configuration](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>For events used in JMS channels Names beginning with <code>_jms</code> or <code>jms</code> (case insensitive) are used only for JMS header properties. You can, however, use properties beginning <code>jms_</code> (case insensitive) for event properties. See Using JMS Header Properties in Incoming and Outgoing Messages on page 93 for more details.</p>
Description	No	Short description of the resource.
Type	No	<ul style="list-style-type: none"> ruleBased: The time event is scheduled by a rule. repeat: Time events are created periodically. Complete the rest of the settings to define the period and number of events per interval.

Field	Global Var?	Description
Interval	No	<p>Used with repeating time events only. Specify the interval between time events using a numerical value and a time unit.</p> <p>TIBCO BusinessEvents creates the first time event immediately after the engine starts and then creates the next time event based on the interval.</p> <p>A repeating time event expires after the completion of the first RTC cycle (that is, the time to live code is set internally to zero).</p> <p>Time units available are milliseconds, seconds, minutes, hours, and days.</p>
Events per Interval	No	<p>Used with repeating time events only. Enter the number of events to create in each Repeat Every interval.</p> <p>Default is 1.</p>

TimeEvent Attributes Reference



When using an event in a rule's form editor, type *eventname@* to see the list of its attributes.

You can use the following attributes in rules to return information about an event instance.

Attribute	Type	Returns
@id	long	The event's unique internal ID.
@closure	String	<ul style="list-style-type: none"> For repeating time events: Not used (Null value). For rule based time events: A string that was specified when the event was scheduled.
@interval	long	<ul style="list-style-type: none"> For repeating time events: The number of milliseconds between creation of successive time events. For rule based time events: Not used (0 value).
@scheduledTime	DateTime	<p>The time scheduled for asserting an instance of this event into the Rete network.</p> <ul style="list-style-type: none"> For repeating time events: calculated based on the time of the last assertion of an instance of this event, and the interval. For rule based time events: specified using the <code>ScheduleTimeEventName()</code> function's <code>delay</code> parameter. See Rule Based TimeEvent Function Reference on page 184.

Attribute	Type	Returns
@ttl	long	<ul style="list-style-type: none">• For repeating time events: always 0 (zero).• For rule based time events: specified using the ontology function's <code>ttl</code> parameter. See Rule Based TimeEvent Function Reference on page 184.

Rule Based TimeEvent Function Reference

Signature Event Schedule*TimeEventName*(long delay, String closure, long ttl)

Domain action

Description For each rule-based time event you create, an ontology function is also created to enable you to schedule the assertion of an instance of the event in a rule action. The function name follows this format: *ScheduleTimeEventName*.

Parameters

Name	Type	Description
delay	long	The event is created (and asserted) the specified number of milliseconds after the rule action executes.
closure	String	The TimeEvent created will simply store the information passed in the closure parameter. you can put the closure string value in rule conditions to specify the time events that will trigger the rule.
ttl	long	The event's time to live. Follows the same rules as the time to live setting in a simple event. See Simple Event Reference on page 167 .

Returns

Type	Description
Event	An instance of the event type specified in the function name (<i>ScheduleTimeEventName</i>).

Using Scheduler Functions (Requires Cache OM)



The TIBCO BusinessEvents Express edition supports only In Memory object management. Therefore it does not support scheduler functions.

You schedule time events as explained in [Overview of Time Events on page 176](#), and you schedule simple events using scheduler functions. Events scheduled using these functions are sent to their default destination.

You might use the event scheduler functions instead of a time event in the following cases:

- If you need to send any event on a schedule, you must use the scheduler. If you schedule a memory-only simple event the schedule is persisted. If you schedule a memory only time event (using the time event's function, not the scheduler functions) then the schedule is not persisted.
- If you need to create many, perhaps thousands, of schedules at one time, performance is better using the event scheduler functions.
- Because schedule management and event sending can be handled by a cache agent, fewer resources in the inference agents are required compared to time events. Additionally, with backing store enabled the schedules are only loaded into memory a batch at a time, reducing the total memory requirement compared to time events.

Task A Create a Scheduler

First you create a scheduler using the function `Cluster.createScheduler()`. You can do this, for example, in a startup rule function, or other rule function or rule:

```
void createScheduler(String schedulerName, long pollInterval, long refreshAhead)
```

- *schedulerName* A unique ID for the scheduler you are creating.
- *pollInterval* TIBCO BusinessEvents checks the scheduler cache every *pollInterval* milliseconds, for scheduler work items whose *scheduledTime* falls in the current interval.
- *refreshAhead* Time in milliseconds (into the future) used to pre-load the scheduled events from the backing store.



In this release, the larger value of *pollInterval* and *refreshAhead* is used for both these settings.

For example:

```
Cluster.createScheduler("myScheduler", 1000, 5000);
```

Task B Schedule the Event to be Sent

You use the function `Cluster.scheduleEvent()` to schedule an event to be sent at a certain time, using a scheduler you created earlier. You can do this in a rule or rule function as needed:

```
void scheduleEvent(String schedulerName, String workKey, SimpleEvent evt, long  
scheduledTime)
```

Where:

- *schedulerName* is the unique ID for the scheduler you created at an earlier time.
- *workKey* is a unique key that identifies the work item (that is, the scheduled task). This key can be used to identify the work item later, for example to cancel it.
- *evt* is the simple event to be scheduled.
- *scheduledTime* is used to specify the time the event is sent to its default destination. The value is interpreted as an absolute time. To schedule a time relative to the present, use the `System.currentTimeMillis()` function, as shown below.

For example:

```
Cluster.scheduleEvent("MyScheduler", myworkKey, Event.createEvent("xslt:// details  
omitted "), System.currentTimeMillis() + 5000)
```

Chapter 11 **Advisory Events**

TIBCO BusinessEvents supports three sorts of events: Simple events (usually referred to just as events); time events, which are timers; and advisory events. This chapter explains how to use advisory events.

Topics

- [Working With Advisory Events, page 188](#)
- [Advisory Event Attributes Reference, page 190](#)

Working With Advisory Events

You never have to add or configure an event of type `AdvisoryEvent`. Advisory events are asserted into the Rete network automatically when certain conditions, for example, exceptions, occur. Add the `AdvisoryEvent` event type to rule declarations to be notified of such conditions. Then use the event attributes in the rule as needed ([Advisory Event Attributes Reference on page 190](#)).

An advisory event expires after the completion of the first RTC cycle (that is, the time to live is set internally to zero).

Uses of Advisory Events

See [Advisory Event Attributes Reference on page 190](#) for details about use of attributes in each of the following uses.

Exceptions in User Code

The TIBCO BusinessEvents engine automatically asserts an advisory event when it catches an exception that originates in user code but that is not caught with the `catch` command of the TIBCO BusinessEvents Exception type.

For information on working with other kinds of exceptions, see [Exception Handling on page 334](#).

Engine is Activated

An advisory event (`engine.primary.activated`) is asserted when an engine has finished starting up and executing startup functions, if any (see Engine Startup and Shutdown Sequence in *TIBCO BusinessEvents Administration*).

Adding an Advisory Event to a Rule

1. Open the rule editor for the rule where you want to use the advisory event. (See [Adding a Rule on page 254](#).)
2. Do one of the following:
 - In the form-based editor, click the plus icon in the Declarations panel and select Advisory Event from the list of resources.
 - In the source editor, add a line like the following in the `declare` block:

```
AdvisoryEvent a;
```

In both cases you can change the alias as desired (from `a` to something else).

3. Save the resource.

The `AdvisoryEvent` event type has no properties. You can use its event attributes in rules to return information about an advisory event. See [Advisory Event Attributes Reference on page 190](#).

Advisory Event Attributes Reference



When using an event in a rule's form editor, type *eventname@* to see the list of its attributes.

The AdvisoryEvent event type has no properties. You can use the following attributes in rules to return information about an advisory event. Add-on products may use additional advisory event types.

Name	Type	Description
@id	long	The event's unique internal ID.
@extId	String	Null.
@category	String	Broad category of advisory: Exception—Used for exceptions not otherwise caught. Engine—Used for engine events. Also used for TIBCO BusinessEvents-ActiveMatrix BusinessWorks integration projects. Deployment —Used for hot deployment
@type	String	Type of advisory within the category. See Table 27, Attributes Used for Each Type of Advisory Event, on page 190 for details.
@message	String	Message content depends on the type of advisory event. See Table 27, Attributes Used for Each Type of Advisory Event, on page 190 for details.

Table 27 Attributes Used for Each Type of Advisory Event

Description	@category	@type	@message
Exceptions arising from user code	Exception	The Java class name of the exception	The stack trace and a message (if the exception includes a message)
TIBCO ActiveMatrix BusinessWorks Activity Failure	Engine	INVOKE BW PROCESS Indicates a failure or cancellation or timeout of the ActiveMatrix BusinessWorks process.	The error message from the failed ActiveMatrix BusinessWorks process, or the timeout message.

Table 27 Attributes Used for Each Type of Advisory Event

Description	@category	@type	@message
Engine is activated	Engine	engine.primary.activated Indicates the engine has been activated.	Engine <i>EngineName</i> activated.
Hot deployment	Deployment	deployment.hotdeploy.success deployment.hotdeploy.fail	Hot deployed project <i>ProjectPath</i> Failed to hot deploy project <i>ProjectPath</i>

Chapter 12 **Concepts**

This chapter explains how to work with concepts.

Topics

- [Overview of Concepts, page 194](#)
- [Adding Concepts and Concept Relationships, page 195](#)
- [Concept Resource Reference, page 197](#)
- [Concept Attributes Reference, page 200](#)

Overview of Concepts

You can add concept definitions so that information that arrives in events or from other sources can be organized and persisted as needed, and used in rules. You can add definitions manually. You can also import database tables as concept definitions. Concept instances are created in rules.

See Also

- To gain an understanding of concepts including a detailed discussion of concept history, and concept relationships, Chapter 3, Concepts in *TIBCO BusinessEvents Architect's Guide*.
- For working in rules, [Working with Concept and Event Properties on page 331](#).
- To configure runtime properties, including backing store behavior, [Chapter 28, Domain Objects Configuration](#).
- For information on using concept dependency and concept model diagrams [Chapter 41, Diagrams, on page 701](#).

Concept Serialization and Handling of Null Value Properties at Runtime

By default, when concept instance objects are serialized to XML, properties with null values are excluded. You can change this behavior so that null values are included. You can also change the XSD for a concept object to allow null values, using the nillable attribute. See [Appendix B, Handling Null Properties, on page 741](#) for details.

Concepts and State Machines

If you are using the TIBCO BusinessEvents Data Modeling add-on product, you can associate a concept with a state machine. See *TIBCO BusinessEvents Data Modeling Developer's Guide* for details.

Adding Concepts and Concept Relationships

This section explains how to work with concepts, and how to set up relationships between concepts.

Adding a Concept

See [Concept Resource Reference on page 197](#) for details on completing values.

To Add a Concept

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the concept and select **New > Concept**. You see the New Concept Wizard.
2. In the Concept Name field, type a name for the concept. In the Description field, type a description.



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. Click **Finish**. You see the Concept Editor.
4. Do the following:
 - Configure the remaining fields as explained in [Wizard and Configuration Tab on page 197](#).
 - As needed, add and configure properties as explained in [Properties on page 198](#).
5. Save the resource.



The ability to add metadata groups and properties is provided for customization purposes and is not documented. Metadata properties are also used for database concepts, a feature available in TIBCO BusinessEvents Data Modeling. See *TIBCO BusinessEvents Data Modeling Developer's Guide* for details.

Adding Concept Relationships

See Concept Relationships in *TIBCO BusinessEvents Architect's Guide* for information about concept relationships.

To Create an Inheritance Relationship

In the child concept resource's Configuration section, identify the parent concept in the **Inherits From** field.

Note that inheritance is set at the concept level. Other relationships are set at the concept property level.

To Create a Containment Relationship

A concept can be contained by only one concept at a time.

1. In the container concept resource's Property section, add and configure a property of type **ContainedConcept**.
2. Configure the property as follows:

Name: Provide a helpful name. For example, "wheels" could indicate that this concept is contained by a "car" concept (in an appropriate project).

Multiple: Check the checkbox if this property is an array.

Policy: Changes Only or All Values.

History: The number of historical values to keep.

To Create a Reference Relationship

A concept can refer to itself, and can be referred to by more than one concept.

1. In the Property tab of one of the concept resources, create a property of type **ConceptReference**.
2. Configure the property as follows:

Name: Provide a helpful name such as `is a lineitem of` to express the relationship.

Multiple: Check the checkbox if this property is an array.

Policy: Changes Only or All Values.

History: The number of historical values to keep.

See **Concept Property History** in *TIBCO BusinessEvents Architect's Guide* for more on this topic.

Editing or Deleting Concept Relationships

You cannot edit or delete the relationship using the concept view graphical user interface. Instead, work with the concept properties directly in the Concept Editor.

Concept Resource Reference



Concept resources are descriptive entities similar to the object-oriented concept of a class. They describe a set of properties. For example, one concept might be Department, and it could include department name, manager, and employee properties.



If you are working with a project imported from a release earlier than 5.0.0, you may see metadata properties. However, do not use them. Instead use the settings and properties in the Domain Objects section of the CDD file as needed. See [Chapter 28, Domain Objects Configuration](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>This field appears on the New Concept Wizard. The name then appears in the title of the concept editor.</p>
Description	No	<p>Short description of the resource.</p> <p>This field appears on both the New Concept Wizard and in the concept editor.</p>
Inherits From	No	<p>If you want this concept to inherit all the properties of another concept, browse to and select that concept.</p> <p>Concepts that are related to each other directly or indirectly by inheritance cannot have distinct properties that share a common name. Therefore the following restrictions apply:</p> <ul style="list-style-type: none"> • If two concepts are related by inheritance, you cannot create a new property in one with a name that is already used in the other. • If two unrelated concepts have properties that share a name, an inheritance relationship cannot exist between them.

Field	Global Var?	Description
State Models	No	State models that are owned by this concept. Models can be added and removed as needed. Requires TIBCO BusinessEvents Data Modeling.
Auto Start State Model	No	If checked then when a concept is asserted, at runtime, its main state machine (if any) is started. If unchecked, then state machines are started using this function in rules: <code>Instance.startStateMachine()</code> The function has a parameter specifying whether child concepts' state machines are also started. For contained concepts, the value of the Auto Start State Machine is not relevant and is not checked. Requires TIBCO BusinessEvents Data Modeling.

Properties

The Properties tab has the following fields:

Field	Global Var?	Description
Name	No	The property name.
Type	No	Any of the following types: <code>String, Integer, Long, Double, Boolean, DateTime, ContainedConcept, ConceptReference</code> When you create a property of type <code>ContainedConcept</code> , you are creating a containment relationship. The concept that you are currently configuring is the container; the concept you specify as a property is the contained concept. When you create a property of type <code>ConceptReference</code> you are creating a property that references another concept. See Concept Relationships in <i>TIBCO BusinessEvents Architect's Guide</i> for more details. Note: For properties of type <code>Double</code> , when a backing store is used, all NaN (Not a Number) values are converted to <code>0.00</code> .

Field	Global Var?	Description
Multiple	No	<p>Check the Multiple checkbox if this property is an array.</p> <p>Consider, for example, an Order concept: In most cases, an Order concept would allow only one value for the customer property but multiple values for the <code>line_item</code> property. Selecting the Multiple checkbox creates a property array.</p>
Policy	No	<p>TIBCO BusinessEvents can record historical values using either of these policies:</p> <p>Changes Only TIBCO BusinessEvents records the value of the property every time it changes to a new value.</p> <p>All Values TIBCO BusinessEvents records the value of the property every time an action sets the value even if the new value is the same as the old value.</p>
History	No	<p>Determines if historical values are stored, and if so how many.</p> <p>The default maximum value is 32767.</p> <p>Zero (0): TIBCO BusinessEvents does not store historical values for the concept. It stores the value without a time and date stamp</p> <p>One or more (>0): TIBCO BusinessEvents stores the property value when the property changes, along with a date and timestamp, up to the number specified. When the maximum history size is reached, the oldest values are discarded as new values are recorded. See Concept Property History in <i>TIBCO BusinessEvents Architect's Guide</i> for more details.</p> <p>Note: Use of a temporal function with a concept that has a history size of 0 may cause a runtime exception.</p>

Metadata

The Metadata section is used in a special way with database concepts. Requires TIBCO BusinessEvents Data Modeling. See *TIBCO BusinessEvents Data Modeling Developer's Guide* for details. It is not used otherwise (except for customization, which is not documented).

Concept Attributes Reference



When using a concept in a rule's form editor, type `conceptname.@` to see the list of its attributes.

You can use the following attributes in rules to return information about a concept instance.

Entity	Attributes	Type	Returns
Concept	@id	long	The concept instance's unique internal ID.
	@extId	string	The concept instance's unique external ID. Optional. The value of the <code>extId</code> is set at creation time (for example, using the <code>Instance.createInstance()</code> function) and cannot be changed after that. Note: The <code>extId</code> value (if set) must be unique across all objects in the cluster. Tip: You can use the property <code>Agent.AgentClassName.checkDuplicates</code> to check for duplicate <code>extIds</code> across the cluster. See CDD Agent Classes Tab Properties Reference on page 501.
	@parent	concept	The parent concept instance. (This is treated as a concept reference in the language.)
ContainedConcept	@id	long	The contained concept instance's unique internal ID.
	@extId	string	The contained concept instance's unique external ID.
	@parent	concept	The parent concept instance. (This is treated as a concept reference in the language.)

Entity	Attributes	Type	Returns
ConceptReference	@id	long	The contained concept instance's unique internal ID.
	@extId	string	The contained concept instance's unique external ID.
	@isSet	boolean	Available for any concept property (not for referenced concepts). For example, you can use: <code>acc.Balance@isSet</code> . It indicates if a property is assigned or not. If a property is assigned, then <code>@isSet</code> returns <code>true</code> .

Chapter 13 **Scorecards**

This brief chapter explains how to work with scorecards.

Topics

- [Understanding and Working With Scorecards, page 204](#)

Understanding and Working With Scorecards

A scorecard is a special type of concept. A scorecard serves as a set of static variables that is available throughout the project. You can use a scorecard resource to track key performance indicators or any other information.

Unlike concepts and events, each scorecard resource is itself a single instance — it is not a description for creation of instances. You create the scorecard at design time. Its values can be viewed and updated using rules.

It is more accurate to say there is one instance of a scorecard per inference agent. Each inference agent in an application has its own instance of the score card. Scorecards are not shared between agents.

Any agent that uses scorecards, and also uses Cache OM, must be assigned a unique key so that the correct scorecard can be retrieved from the cache. The key is set in the Processing Unit tab of the CDD. See [Configuring Processing Units \(All OM Types\) on page 513](#).

It is not necessary to add scorecards to the declaration of a rule. Because there is only one instance of each scorecard in a deployed TIBCO BusinessEvents agent, any change causes all rules that use the scorecard in their conditions to be evaluated.



The `Instance.isModified()` function works differently with scorecards than with concepts. There is only one instance of a scorecard per agent, rather than one per RTC. So after a scorecard is modified it will return true until the agent is restarted.

(In the case of a concept instance, `Instance.isModified()` returns true after the instance has been modified only for the rest of the RTC in which it is modified.)

See Also

- [Chapter 12, Concepts, on page 193](#).
- [Chapter 41, Diagrams, on page 701](#) for information on using scorecard dependency diagrams.

Adding a Scorecard

Configuring a scorecard is similar to configuring a concept, except that scorecards do not have relationships.

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the scorecard and select **New > Scorecard**. You see the New Scorecard Wizard.
2. In the Scorecard Name field, type a name for the scorecard. In the Description field, type a description.



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. Click **Finish**. You see the Scorecard Editor.
4. Add and configure properties as explained in [Concept Resource Reference on page 197](#).
5. Save the resource.



The ability to add metadata groups and properties is provided for customization purposes and is not documented.

Using a Scorecard in Rules

After configuring a scorecard resource, use rules to gather the information you need in the scorecard. To access the scorecard in a rule, use this syntax:

```
folder . folder . scorecard . property
```

For Example:

```
int i = SalesFolder.StatsScorecard.numOrdersProperty;
```

Scorecard Resource Reference



Scorecard configuration is the same as concept configuration, except that scorecards have no relationships with each other. Scorecards, therefore, have none of the properties used for setting up relationships. Scorecard properties can be of any primitive type.

See [Concept Resource Reference on page 197](#) for details about scorecard properties

Chapter 14 **Domain Models**

Domain model resources enable you to control user input for decision tables and test data. This chapter describes how to add and import domain models and associate them with entity properties.

Topics

- [Overview of Domain Models, page 208](#)
- [Adding a Domain Model, page 209](#)
- [Importing and Exporting Domain Models, page 212](#)
- [Associating Domain Models with Properties, page 215](#)
- [Validating Data in Domain Models, page 216](#)

Overview of Domain Models

You can add domain models for concept, event, and scorecard properties. A domain model specifies the values that you may find useful for defining ontology item properties. For example, instead of typing text for a certain concept property, you can pick a value from a list, or enter a value within a predefined range.

A domain model can extend another domain model. When defining a domain model you can specify which domain model it inherits from.

Domain model entries are case sensitive. For example, m and M are recognized as different entries.

Domain models can be used in decision tables and in test data for Tester. Decision tables are available in the TIBCO BusinessEvents Decision Manager add-on. See TIBCO BusinessEvents Decision Manager documentation for details about working with decision tables.



You cannot add domain models for concept properties of the `ContainedConcept` or `ConceptReference` type.

Setting up domain models has two steps:

1. Add domain models. For example, create a folder and add all the models you need for a project in this folder. See [Adding a Domain Model on page 209](#) for details.

You can also import (and export) domain models. See [Importing and Exporting Domain Models on page 212](#).

2. Associate domain models with properties. See [Associating Domain Models with Properties on page 215](#) for details.

When creating a decision table, domain models control the values you can use for a given property.

Domain Model Value Descriptions for User Friendly Presentation

All domain model values can have optional descriptions that appear in the domain model editor. A preference determines whether domain model values or their descriptions appear in decision table cells. For some applications displaying descriptions can make the table easier for users to understand. For example suppose the value is a code such as 23, and the description is North West. Users will find it is easier to work with the description than the code. As another example, for a Boolean data type, the description can provide words such as Accepted and Rejected for the values True or False.

Adding a Domain Model

You can store domain models as desired, for example, in a folder called `DomainModels`. For each domain model, you create a set of domain entries, where each entry represents a valid value for the entity property that uses the domain model.

After you add a domain model, associate it with a property. See [Associating Domain Models with Properties on page 215](#)

To Add a Domain Model

1. Right-click the folder where you want to store the domain model, and select **New > Domain Model**. You see the New Domain Model Wizard.

Alternatively, right-click a resource, and select **New > Other**. In the New dialog, select **Domain Model** under TIBCO BusinessEvents.

2. In the Domain Model Name field, type a name for the domain model. In the Description field, type a description.



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. Click **Finish**. You see the Domain Model editor.
4. In the Domain Type field, select the data type for this domain model. See [Supported Data Types on page 210](#) for supported types.
5. As needed complete the Inherits From field. If this domain model extends another domain model, browse to and select that domain model.
6. In the Domain Entries section, click **Add**. An empty row appears in the table of entries.

You can also select rows and click **Duplicate** to duplicate (and then modify) selected rows.

You can select a row and click **Remove** to remove individual rows that are not needed.
7. For each row you add, enter a description (optional) and in the Details section, define the domain model entry.

The Details section presents appropriate fields for defining the type of domain model entry you selected in [step 4](#). See [To Add Domain Entries on page 210](#) for examples.

- When you have created the entries for the domain model, save the domain model resource.

To Add Domain Entries

When you add a domain model, you first select its data type. The Domain Model editor then displays an appropriate user interface for defining domain model entries of that data type.

The Domain Entries section is the same for all types. The Details section changes depending on what data type you selected for the domain model.

Supported Data Types

Domain Models support the following data types:

- String
- Integer
- Long
- Double
- Boolean
- DateTime

Sections below show the user interface for each type.

String

String entries are simple text strings.

Integer, Double,



Numeric values in a String domain type When you use a numeric value (Integer, Double, or Long) in a domain model of type String, TIBCO BusinessEvents adds double quotes around the value. (These are visible after you save and reopen the domain model.)

Long

In a single domain model, you can enter single values, range values, or a mixture of both. Acceptable values for integer, long, and double domain entries are the same as for the equivalent Java datatypes.



Text values in a numeric data type When you enter text in a domain model of type Integer, Double or Long, TIBCO BusinessEvents marks it in red color, and indicates you to correct it.

The user interface is similar for all numeric datatypes.

Note that you define whether each end of the range is inclusive. For complete coverage, ensure that there is no gap and no overlap between ranges through consistent use of the Included check box. For example, the figure shows that the lower bound of the Senior range, 55, is not included.

Boolean Boolean entry values are always `true` or `false`. The description can give the meaning of the pair of choices, such as male or female, supported or unsupported, eligible or ineligible and so on.

DateTime In a single domain model, you can enter single values, range values, or a mixture of both. You specify a date and a time. A date and time picker enable you to select the date and time:

If you do not want to specify a time of day, set the time to midnight (12:00 AM) for the start date, and to a minute before midnight (11:59:59 PM) for the end date of an inclusive range, or for a single date.

Importing and Exporting Domain Models

You can import domain model information from a database, from a Microsoft Excel spreadsheet, and from the source of a database concept property. (Database concepts are available in the TIBCO BusinessEvents Data Modeling add-on).

After you import a domain model, associate it with a property. See [Associating Domain Models with Properties on page 215](#)

Import From the Source of a Database Concept

You can import all values from the database column that corresponds to a database concept property for use as the domain model entry for that property. This is explained in *TIBCO BusinessEvents Data Modeling Developer's Guide*.

Export To and Import From Excel

You can export a domain model to an Excel spreadsheet. The description for each domain model entry goes in column A, and the value goes in column B. When you import from an Excel spreadsheet, similarly, column A is used as the description, and column B is used for the value.

In TIBCO BusinessEvents Decision Manager, you can export decision tables to Excel and import them from Excel. When you do so, you also export and import domain models associated with properties used in a decision table column. See *TIBCO BusinessEvents Decision Manager User's Guide* for details.

To Import Domain Model Entries from Excel

1. In BusinessEvents Studio Explorer, do one of the following:
 - Right-click the folder where you want to create the domain model and select **Import > TIBCO BusinessEvents > Domain Model**.
 - In BusinessEvents Studio Explorer, select any item in the project entity and select **File > Import > TIBCO BusinessEvents > Domain Model**.
2. Click **Next**. You see the Import Domain Model Wizard.

If you invoked the import wizard by right-clicking a folder, that folder is selected as the parent folder. You can choose a different one as desired.
3. In the Domain Import Source field, select **EXCEL**.
4. In the File Name field, enter a name for the domain model resource. Optionally enter a description.

5. In the Data Type field, select the appropriate data type for the domain model and click **Next**. You see the next wizard page.
6. In the Select Excel File to Import field, browse and select the Excel file that has the domain model information.
7. Click **Finish**. You see a message, "Domain Import Successful." Click **OK**.
8. You see the Domain Model editor. The column values appear as domain entries. You can add, edit, duplicate, and remove entries as appropriate.

To Export Domain Values to Excel

1. In BusinessEvents Studio Explorer do one of the following:
 - Right-click the domain you want to export and select **Export > TIBCO BusinessEvents > Export Domain to Excel**.
 - In BusinessEvents Studio Explorer, select any item in the project entity and select **File > Export > TIBCO BusinessEvents > Export Domain to Excel**.
2. Click **Next**. You see the Export to Microsoft Excel File Wizard.
3. In the Select Excel File to Export to field, browse and select the Excel file to which you want to export the domain model information.

To create a new Excel file, specify a filename that does not yet exist. If you specify an existing Excel file, the file contents are replaced with the exported domain model information.
4. In the Select Domain to Export area, expand the display and select the domain model you want to export.

If you invoked the export wizard by right-clicking a domain model, that domain model is selected. You can choose a different one as desired.
5. Click **Finish**. You see a message, "Export Domain Model to File Successful." Click **OK**.

Import From a Database Table

When you import domain model information from a database, the result set from the SQL query is transformed for use as domain model entries. This feature is supported with Oracle Database.

To Import Domain Model Entries from a Database Table

1. Add a JDBC Connection resource and configure it to connect to the database from which you want to import the domain model. See [JDBC Connection on page 230](#) for details.

2. In BusinessEvents Studio Explorer, do one of the following:
 - Right-click the folder where you want to create the domain model and select **Import > TIBCO BusinessEvents > Domain Model**.
 - In BusinessEvents Studio Explorer, select any item in the project entity and select **File > Import > TIBCO BusinessEvents > Domain Model**.
3. Click **Next**. You see the Import Domain Model Wizard.

If you invoked the import wizard by right-clicking a folder, that folder is selected as the parent folder. You can choose a different one as desired.
4. In the Domain Import Source field, select **DATABASE_TABLE**.
5. In the File Name field, enter a name for the domain model resource. Optionally enter a description.
6. In the Data Type field, select the appropriate data type for the domain model and click **Next**. You see the next wizard page.
7. In the JDBC Resource URI field, browse and select the JDBC Connection resource that connects to the database you want to use.

The connection information from the JDBC Connection resource displays. You can override it here as desired.

Click **Next**. You see a list of tables in the database. Expand the list of tables to see the columns that match the datatype specified in [step 6](#).
8. Do one of the following:
 - Select one or more columns, then click the **Create Domain for Selected Columns** button. Values of all columns are used for the domain model entries.
 - Click **Advanced** and enter an SQL query whose resultset is used to create the domain entries. This option enables you to make use of joins, where clauses, and so on. Then click **Execute Query**.
9. Click **Finish**. You see a message, "Domain Import Successful." Click **OK**.
10. You see the Domain Model editor. The column values appear as domain entries. You can add, edit, duplicate, and remove entries as appropriate.

Associating Domain Models with Properties

You can associate domain models with concept, event, and scorecard properties.

If domain models have the same data type as that of the properties, you can associate multiple domain models with multiple properties.

For example, an event and a concept have an `OrderID` property defined the same way, and both the `OrderID` properties use an `OrderIDModel` domain model. In this case, the `OrderIDModel` domain model is associated with the `OrderID` property of the event, as well as of the concept. However, the concept or the event property is associated only with the `OrderIDModel` domain model.

To Associate Domain Models with a Property

1. Do one of the following:
 - In BusinessEvents Studio Explorer, expand the concept, event, or scorecard to display its properties. Right-click the desired property and select **Associate Domains**.
 - In BusinessEvents Studio Explorer, open the desired concept, event, or scorecard editor. In the Domain Model field of the desired property, click the browse icon.

The Associate Domains dialog appears.

2. Expand the project tree to display domain models, and select one or more as appropriate.

The displayed domain models have the same data type as that of the property.
3. Save the domain model resource.

Validating Data in Domain Models

You can validate duplicate entries and a mismatch of upper and lower range values while defining domain models.

Duplicate Domain Values Not Allowed

Each domain value must be unique. If you accidentally enter duplicate values, the Problems view displays helpful information.

The Resource column displays the duplicate value.

Mismatching Range Values Not Allowed

If you accidentally enter mismatching values in the Lower and Upper fields for range values, the Problems view displays helpful information.

Chapter 15 Shared Resources

This chapter describes a set of resources that can be used for various purposes in your projects.



A RuleServiceProvider Configuration shared resource is available In TIBCO BusinessEvents Studio. It can be used during migration to display a migrated version 3.x resource for informational purposes. For integration with TIBCO ActiveMatrix BusinessWorks projects, use the RuleServiceProvider Configuration activity, available in the TIBCO Designer palette of TIBCO BusinessEvents Activities. See [Chapter 38, ActiveMatrix BusinessWorks Integration, on page 637](#).

Topics

- [Adding a Shared Resource, page 218](#)
- [ActiveSpaces Connection, page 220](#)
- [Hawk Connection, page 223](#)
- [HTTP Connection, page 225](#)
- [Identity Resource, page 228](#)
- [JDBC Connection, page 230](#)
- [JMS Application Properties, page 234](#)
- [JMS Connection, page 235](#)
- [JNDI Configuration, page 242](#)
- [Rendezvous Transport, page 245](#)

Adding a Shared Resource

See the section [Working with External Library and Custom Function Paths on page 8](#) for actions you must do to run, test, or debug projects that use third-party libraries in TIBCO BusinessEvents Studio.

If you are adding a JDBC connection for a backing store, see the procedure, [Add a JDBC Connection Resource \(Now or Later\) on page 537](#).



It's a good idea to use global variables in shared resources so that projects can be quickly adapted to run in different environments.

To Add a Shared Resource

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the shared resource and select **New > Other**.
2. In the Select a Wizard dialog, expand TIBCO Shared Resources and select the resource type you want to add (see list in [step 3](#)) and click **Next**.
3. In the new resource wizard enter a name in the File Name field and click **Finish**.

You see the editor for the shared resource you selected. See the following sections for guidelines on completing the fields:

- [ActiveSpaces Connection, page 220](#)
 - [Hawk Connection, page 223](#)
 - [HTTP Connection, page 225](#)
 - [Identity Resource, page 228](#)
 - [JDBC Connection, page 230](#), and also see [Enabling the Test Connection Feature, page 219](#)
 - [JMS Application Properties, page 234](#)
 - [JMS Connection, page 235](#), and also see [Enabling the Test Connection Feature, page 219](#)
 - [JNDI Configuration, page 242](#)
 - [Rendezvous Transport, page 245](#)
4. (Optional) View the Source tab. The source tab shows the XML source format for the shared resource. You can edit the XML source directly, but this is not recommended because of the risk of error. The resource editor provides validation checks.

Enabling the Test Connection Feature

To make the Test Connection feature work for JMS Connection and JDBC Shared Connection shared resources, copy the relevant JAR files to the following directory:

BE_HOME/lib/ext/tpcl

(This directory is referenced in the extended classpath in the file *BE_HOME/studio/eclipse/configuration/studio.tra*.)



If TIBCO BusinessEvents Studio is running when you copy the file or files, you must restart it for the test connection feature to work.

For TIBCO Enterprise Message Service copy the *jms-2.0.jar* and *tibjms.jar* to the above location.

For WebSphere MQ, copy the MQ JAR files and the binding file to the above location.

For DBMS products, copy the supported driver file you are using for the DBMS product.

ActiveSpaces Connection

The ActiveSpaces Connection resource describes the connection to a TIBCO ActiveSpaces metaspace. This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
File Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.</p> <p>The name then appears in the title of the resource.</p>
Configuration		
Description	No	Short description of the resource.
Metaspace Name	Yes	<p>Specifies a TIBCO ActiveSpaces metaspace. A metaspace is an instance of a cluster of application processes deployed using TIBCO ActiveSpaces. The application processes are typically deployed on multiple hosts that are interconnected by a network.</p> <p>The default value for this field is 'ms'.</p>
MemberName	Yes	Name the AS channel will take when connecting to ActiveSpaces metaspace

Field	Global Var?	Description (Cont'd)
Multicast URL	Yes	<p>Specifies the URL to be used for multicast discovery. The multicast URL is specified using the following format:</p> <pre>tibpgm://[dport]/[interface]/[discovery IP multicast address]/[option=value;]*</pre> <p>where,</p> <ul style="list-style-type: none"> • <i>dport</i> specifies the destination port used by the PGM transport. If not specified, it will use the default value of 7888. • <i>interface;discovery IP multicast address</i> specifies the address of the interface to be used for sending discovery packets, and the discovery group address to be used. If not specified, it will default to the default interface and discovery address, 239.8.8.8. • <i>option=value</i> specifies a semicolon separated list of optional Pragmatic General Multicast (PGM) transport arguments. For example: <code>source_max_trans_rate=100000000</code> (in bits per second) would limit the PGM transport to limit its transmission rate to 100 megabits per second.
Unicast Url	Yes	<p>Specifies the URL to be used for unicast discovery. The URL specifies the interface and port to be used to create the listening socket on.</p> <p>The unicast URL is specified using the following format:</p> <pre>tcp://ip1[:port1];ip2[:port2],...</pre> <p>where any number of <i>ip[:port]</i> well known address can be listed. If port is not specified, the default port number value of 50000 will be assumed.</p>
EnableSecurity	No	Enable security.

Field	Global Var?	Description (Cont'd)
SecurityRole	No	<p>Indicate the security role of the AS channel in the metaspace: Controller or Requestor:</p> <ul style="list-style-type: none"> • Controller is dedicated to enforcing a security domain's defined security behavior for a metaspace associated with the security domain. Security domain controllers are the only discovery nodes in a metaspace. • Requestor just requires access to the data in the data grid, such as a seeder or a leech, and which need to be authorized by a controller. Requesters can never be used a discovery nodes.
PolicyFile	Yes	Policy file contains the security settings that the Controller enforces. It is generated via the as-admin utility.
TokenFile	Yes	Token file is used by Requestor to connect to a metaspace whose security is defined in the policy file.

Hawk Connection

The Hawk Connection resource describes the connection to a TIBCO Hawk domain through a specific transport. This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
File Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 . Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name. The name then appears in the title of the resource.
Configuration		
Description	No	Short description of the resource.
Transport	No	Type of transport used by Hawk domain: Rendezvous or EMS.
Domain	Yes	Name of the Hawk domain.
Transport: Rendezvous		
RV_Service	Yes	Specifies the User Datagram Protocol (UDP) service group used by the TIBCO Rendezvous daemon for session communications. A service can be specified either by its name or its port number. The default configuration uses the service port number 7474.

Field	Global Var?	Description (Cont'd)
RV_Network	Yes	Specifies the network to be used for outbound session communications when a computer is connected to more than one network. A network can be specified by its name or by its IP address.
RV_Daemon	Yes	<p>Specifies the TIBCO Rendezvous daemon that will handle communication for the session. A local daemon is specified by the communications type (always tcp) and a socket number (e.g., 7474). The default configuration uses the local daemon with the TCP socket number 7474.</p> <p>Specify a remote daemon by inserting its host name or IP address between the tcp entry and the port number of the daemon parameter, for example:</p> <pre>tcp:remote_computer:7800</pre>
Transport: EMS		
EMS_ServerURL	Yes	<p>Specifies the EMS server to connect to. The server URL is typically provided in the following format:</p> <pre>protocol://hostname:port-number</pre> <p>For example, tcp://myhost:7222</p>
EMS_UserName	Yes	Specifies the username used to connect to the EMS server.
EMS_Password	Yes	Specifies the password for the username used to connect to the EMS server.

HTTP Connection

The HTTP Connection resource describes the characteristics of the connection used to receive incoming HTTP requests. This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

The HTTP Connection resource can specify that the HTTPS (secure sockets layer or SSL) protocol must be used by clients. If this is enabled, you can configure the SSL parameters for the HTTP server using the Configure SSL Button. See [Configure SSL on page 226](#) for more information.



If you have multiple HTTP Connection resources specified by multiple HTTP Receiver process starters, the HTTP servers require that all of the connections must be valid to initialize all HTTP Receivers. Therefore, make certain that all HTTP Connection resources have valid configurations before testing or deploying the project.

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
File Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.</p> <p>This field appears on the New HTTP Connection Wizard. The name then appears in the title of the resource.</p>
Configuration		
Description	Yes	Short description of the resource.

Field	Global Var?	Description (Cont'd)
Host	Yes	<p>Specifies the name of the host that accepts the incoming requests. For machines that have only one network card, the default value <code>localhost</code> specifies the current machine. For machines that have more than one network card, this field specifies the host name of the card that will be used to accept incoming HTTP requests.</p> <p>Note Only one HTTP server can be started on each port. Therefore, if you have a machine with multiple network cards, make certain that all HTTP Connection resources that use the same host name specify different port numbers.</p> <p>Note If there is more than one network card on the machine, and you specify <code>localhost</code> in this field, then all network cards on the machine will listen for incoming HTTP requests on the specified port.</p>
Port	Yes	<p>Port number on which to listen for incoming HTTP requests.</p>
Use SSL	No	<p>Specifies whether incoming requests must use the HTTPS (secure socket layer, or SSL) protocol. This protocol authenticates the server to the client, and optionally authenticates the client to the server.</p> <p>Enabling this field allows you to specify SSL parameters with the Configure SSL button (see Configure SSL on page 226).</p>

Configure SSL

The SSL Configuration for HTTPS Connections dialog (accessed by configuring the Configure SSL button) allows you to specify the SSL parameters for the HTTP connection.

The following are the fields in the SSL Configuration for HTTPS Connections dialog:



The HTTPComponent server type does not support Entrust based SSL.

Field	Description
Requires Client Authentication	<p>Checking this field requires clients to present their digital certificate before connecting to the HTTP server.</p> <p>When this field is checked, the Trusted Certificates Folder becomes enabled so that you can specify a location containing the list of trusted certificate authorities.</p>
Trusted Certificates Folder	<p>This field is only applicable when the Requires Client Authentication field is checked.</p> <p>This field specifies a folder in the project containing one or more certificates from trusted certificate authorities. This folder is checked when a client connects to ensure that the client is trusted. This prevents connections from rogue clients.</p>
Identity	<p>This is an Identity resource that contains the HTTP server's digital certificate and private key. See Identity Resource on page 228 for more information.</p>
Strong Cipher Suites Only	<p>When checked, this field specifies that the minimum strength of the cipher suites used can be specified with the <code>bw.plugin.security.strongcipher.minstrength</code> custom engine property. See <i>TIBCO ActiveMatrix BusinessWorks Administration</i> for more information about this property. The default value of the property disables cipher suites with an effective key length below 128 bits.</p> <p>When this field is unchecked, only cipher suites with an effective key length of up to 128 bits can be used.</p>

Identity Resource

The Identity resource encapsulates information that may be used to authorize a user, connection, and so forth. The information you supply changes depending on the type of Identity resource you want to use. This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

The identity certificate location, its type, and password can be specified as global variables.

Wizard and Configuration Tab

The New Identity Resource Wizard and the Configuration tab of the Identity Resource have the following fields.

Field	Global Var?	Description
Wizard		
File Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.</p> <p>This field appears in the New Identity Resource Wizard. The name then appears in the title of the resource.</p>
General		
Description	Yes	Short description of the resource.
Type	No	The type of identity resource: Identify File, Certificate/Private Key, or Username/Password (the default). See sections below for details.
Identity File		
Use this option if the certificate includes the private key information in the same file.		
URL	Yes	Location of the certificate (which includes the private key).

Field	Global Var?	Description (Cont'd)
File Type	No	Choose the certificate file type from the drop-down list: Entrust JCEKS JKS PEM PKCS12
Password	Yes	Password for the certificate.
Certificate/Private Key Identity		
Use this option if the private key and the certificate are in two separate files.		
Certificate URL	No	Location of the certificate. Click the browse icon or type in a URL.
Key URL	No	Location of the private key file associated with the certificate.
Key Password	Yes	Password used for private key.
Username/Password		
Use this option if you want to use a username and password for authentication and do not want to use a certificate.		
Username	Yes	Name of the user for this identity.
Password	Yes	Password for the user for this identity.

JDBC Connection

The JDBC Connection resource describes a JDBC connection. JDBC connections are used with backing stores (see [Chapter 30, JDBC Backing Store Setup, on page 523](#)) and with database concepts (available in the TIBCO BusinessEvents Data Modeling add-on). This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
File Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 . Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.
Configuration		
Description	Yes	Short description of the resource.
Connection Type	No	Specifies the kind of JDBC connection you wish to create. The connection type can be one of the following: <ul style="list-style-type: none"> • JDBC • JNDI The type of connection determines the other configuration fields that appear.

Field	Global Var?	Description (Cont'd)
JDBC Connection Type Configuration Fields		
TIP: Using global variables makes the project more portable		
JDBC Driver	Yes	<p>The name of the JDBC driver class. You can select from a list of drivers or enter a driver manually. Listed drivers are as follows:</p> <pre>oracle.jdbc.OracleDriver (thin) com.ibm.db2.jcc.DB2Driver (supported for database concepts only) com.microsoft.sqlserver.jdbc.SQLServerDriver com.mysql.jdbc.Driver com.sybase.jdbc3.jdbc.SybDriver</pre> <p>When you select a driver, the Database URL field is populated with a template for the URL for the driver.</p>
Database URL	Yes	The URL to use to connect to the database. A template of the URL is supplied for the selected JDBC driver. You must supply the portions of the URL that are in angle brackets, for example, the host, port number, and database instance name.
Maximum Connections	Yes	<p>The maximum number of database connections to allocate. The default maximum is 10. The minimum value that can be specified is 1.</p> <p>See Connection Pooling on page 233 for more details, including related settings that override this setting.</p>
User Name	Yes	User name to use when connecting to the database.
Password	Yes	Password to use when connecting to the database.
Login Timeout	Yes	Time (in seconds) to wait for a successful database connection. Only JDBC drivers that support connection timeouts can use this configuration field. If the JDBC driver does not support connection timeouts, the value of this field is ignored. Most JDBC drivers support connection timeouts.

Field	Global Var?	Description (Cont'd)
JNDI Connection Type Configuration Fields		
JNDI DataSource Name	Yes	The JNDI name specified for the DataSource.
Use Shared JNDI Configuration	No	When this field is checked, the JNDI Configuration field appears, allowing you to choose a shared JNDI Configuration resource. When this checkbox is unchecked, configuration fields appear.
JNDI Configuration	No	This field only appears when the Use Shared JNDI Configuration field is checked. This field allows you to choose a JNDI Configuration shared resource that specifies the JNDI connection information.
JNDI Context Factory	No	The initial context factory class for accessing JNDI. (<code>javax.naming.Context.INITIAL_CONTEXT_FACTORY</code>). You can choose from the drop down list of supported classes, or you can type in a different <code>InitialContextFactory</code> class name.
JNDI Context URL	Yes	The URL to the JNDI service provider (<code>javax.naming.Context.PROVIDER_URL</code>). An example URL is provided when one of the supported JNDI context factory classes is selected. See your JNDI provider documentation for the syntax of the URL.
JNDI User Name	Yes	The user name to use when logging into the JNDI server (<code>javax.naming.Context.SECURITY_PRINCIPAL</code>). If the JNDI provider does not require access control, this field can be empty.
JNDI Password	Yes	The password for logging into the JNDI server (<code>javax.naming.Context.SECURITY_CREDENTIALS</code>). If the JNDI provider does not require access control, this field can be empty.

Connection Pooling

TIBCO BusinessEvents creates a pool of JDBC connections for every JDBC Connection shared resource that uses the JDBC connection type. The maximum size of this pool is specified by the Maximum Connections configuration field.

Resources such as backing stores and database concepts that use this JDBC Connection resource are given a connection from the pool. Once the maximum number of connections is reached, resources requesting a connection cannot proceed. Once a connection is freed by an activity, the connection is returned to the pool. Connections that are left open will eventually time out and be closed. These connections can be reopened at a later time, until the maximum number of connections specified in this field is reached.

For backing store connections, you can use additional connection pool properties, which override equivalent settings in the JDBC Connection resource. See [Configuring Backing Store Settings and Properties on page 456](#) for details.

Test Connection Button

The Test Connection button allows you to test the connection specified in the configuration of this resource. See [Enabling the Test Connection Feature on page 219](#) for a step you must take to enable the connection to work.

JMS Application Properties

The JMS Application Properties resource describes any JMS message properties that a JMS application expects. These properties can then be added.

For procedures see [Adding a Shared Resource on page 218](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
File Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 . Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.
Description	Yes	Short description of the resource.
Configuration		
Properties Table A table listing any application-specific properties. Use the + and x buttons to the right of the table to add and delete properties. Use the up and down arrow buttons to move selected properties to the desired location in the table.		
Property Name	Yes	Name of the column.
Type	Yes	The datatype of the property. Double-click the cell to cause a drop down list of valid JMS datatypes to appear, and choose a value.
Cardinality	No	Specifies whether the property is optional or required. Double-click the cell to cause a drop down list of two values to appear and select optional or required.

JMS Connection

JMS Connection resource describes a JMS connection. This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
File Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.</p> <p>This field appears on the New JMS Connection Wizard. The name then appears in the title of the resource.</p>
Configuration Tab		
Description	Yes	Short description of the resource.
User Name	Yes	<p>User name to use when logging into the JMS server.</p> <p>If the JMS provider does not require access control, this field can be empty.</p> <p>Not all JMS servers require user names and passwords. Refer to your JMS provider documentation and consult your system administrator to determine if your JMS server requires a user name and password.</p>
Password	Yes	<p>Password to use when logging into the JMS server.</p> <p>If the JMS provider does not require access control, this field can be empty.</p>

Field	Global Var?	Description (Cont'd)
Auto-Generate Client ID	No	<p>Checking this field specifies you wish to automatically generate the client ID if no client ID is specified or if the specified ID is already in use. When this field is checked, if a value is specified in the Client ID field, an attempt is made to use the specified value. However, if the specified value is already in use, a new client ID is generated.</p> <p>If this field is not checked, then the value specified in the Client ID field is used. If no value is specified in the Client ID field, then no client ID is set. If the value specified in the Client ID field is already in use, an exception is thrown.</p>
Client ID	Yes	<p>Client ID for the connection. Typically JMS providers have a provider-specific format for client IDs. See your JMS provider's documentation for more information about client IDs. Each connection must use a unique Client ID.</p> <p>You cannot use the same JMS Connection resource for accessing both topics and queues. You should create separate JMS Connection resources if you wish to access both topic and queue destinations.</p>
Use SSL	No	<p>Specifies whether you wish to use SSL for the connection to the JMS server. SSL is used when the Use SSL checkbox is checked. Click the Configure SSL button to configure the SSL connection parameters.</p> <p>Note: SSL is supported only when using TIBCO Enterprise Message Service.</p> <p>See Configure SSL on page 239 for more information.</p>
Use JNDI for Connection Factory	No	<p>Specifies whether JNDI should be used to look up the ConnectionFactory object. If this field is unchecked, the Provider URL and Use XA Connection Factory fields appear. If this field is checked, JNDI configuration fields appear.</p>
Provider URL	Yes	<p>This field is only available when the Use JNDI for Connection Factory field is unchecked.</p> <p>This is the URL to use to connect to the JMS server.</p>
Use XA Connection Factory	No	<p>This field is only available when the Use JNDI for Connection Factory field is unchecked.</p> <p>When checked, this field specifies that an XA connection factory is to be used.</p>

Field	Global Var?	Description (Cont'd)
Connection Factory SSL Password	Yes	<p>This field is only available when the Use SSL checkbox is checked, and the User Shared JNDI Configuration checkbox is unchecked.</p> <p>The SSL configuration is specified in the ConnectionFactory object, except for the client SSL password.</p> <p>You can specify your client SSL password in this field, or you can leave this field empty. If your password is not specified, the private key password is used.</p>
Use Shared JNDI Configuration	No	<p>When this field is checked, the JNDI Configuration field appears. It allows you to choose a JNDI Configuration.</p> <p>When this field is unchecked, JNDI configuration fields appear.</p>
JNDI Configuration	No	<p>This field only appears when the Use Shared JNDI Configuration field is checked.</p> <p>This field allows you to choose a JNDI Configuration that specifies the JNDI connection information. See JNDI Configuration on page 242 for more information.</p>
JNDI Context Factory	Yes	<p>This field only appears when the Use Shared JNDI Configuration field is unchecked. Required. The initial context factory class for accessing JNDI. (<code>javax.naming.Context.INITIAL_CONTEXT_FACTORY</code>). You can choose from the drop down list of supported classes or you can type in a different InitialContextFactory class name.</p> <p>Note: TIBCO BusinessEvents attempts to find the class. However, you may need to add the Java file supplied by your JNDI service provider to the CLASSPATH environment variable to use JNDI.</p>
JNDI Context URL	Yes	<p>This field only appears when the Use Shared JNDI Configuration field is unchecked. Required. This is the URL to the JNDI service provider (<code>javax.naming.Context.PROVIDER_URL</code>). An example URL is provided when one of the supported JNDI context factory classes is selected.</p> <p>See your JNDI provider documentation for the syntax of the URL.</p>

Field	Global Var?	Description (Cont'd)
JNDI User Name	Yes	This field only appears when the Use Shared JNDI Configuration field is unchecked. User name to use when logging into the JNDI server (<code>javax.naming.Context.SECURITY_PRINCIPAL</code>). If the JNDI provider does not require access control, this field can be empty.
JNDI Password	Yes	This field only appears when the Use Shared JNDI Configuration field is unchecked. Password to use when logging into the JNDI server (<code>javax.naming.Context.SECURITY_CREDENTIALS</code>). If the JNDI provider does not require access control, this field can be empty.

Test Connection Button

The Test Connection button allows you to test the connection specified in the configuration of this resource. See [Enabling the Test Connection Feature on page 219](#) for a step you must take to enable the connection to work.

When JNDI is used (that is, when the Use JNDI for Connection Factory checkbox is checked), the Test Connection button works only when the JNDI-related fields on the Configuration and Advanced tab are correctly specified.

Advanced Tab

The Advanced tab has the following fields.

Field	Global Var?	Description
Topic Connection Factory	Yes	<p>This field is only available when the Use JNDI for Connection Factory field on the Configuration tab is checked.</p> <p>The <code>TopicConnectionFactory</code> object stored in JNDI. This object is used to create a topic connection with a JMS application.</p> <p>See your JNDI provider documentation for more information about creating and storing <code>TopicConnectionFactory</code> objects.</p>

Field	Global Var?	Description (Cont'd)
Queue Connection Factory	Yes	<p>This field is only available when the Use JNDI for Connection Factory field on the Configuration tab is checked.</p> <p>The <code>QueueConnectionFactory</code> object stored in JNDI. This object is used to create a queue connection with a JMS application.</p> <p>See your JNDI provider documentation for more information about creating and storing <code>QueueConnectionFactory</code> objects.</p>
Optional JNDI Properties	No	<p>Any additional properties to supply for the connection. You specify a name, datatype, and value for each property.</p> <p>These properties are typically vendor-specific. See your JNDI provider documentation for more information about the available properties.</p>

Configure SSL

The SSL Configuration button allows you to configure the SSL connection parameters.



When using JNDI to lookup the JMS Connection factory, the parameters `ssl_identity` and `ssl_verify_host` must be specified in the `factories.conf` file of the Enterprise Message Service server.

The following table describes the SSL Configuration dialog.

Field	Description
Trusted Certificates Folder	<p>Location of the trusted certificates on this machine. The trusted certificates are a collection of certificates from servers to whom you will establish connections. If the server you wish to establish a connection presents a certificate that does not match one of your trusted certificates, the connection is refused.</p> <p>This prevents connections to unauthorized servers.</p> <p>Trusted certificates must be imported into a folder, and then you can select the folder in this field.</p>

Field	Description (Cont'd)
Identity	<p>The location of the client certificate.</p> <p>You only need to specify the client certificate when the JMS server requires client authentication.</p> <p>See Identity Resource on page 228 for more information.</p>
Trace	<p>Specifies whether SSL tracing should be enabled during the connection. If checked, the SSL connection messages are logged and sent to the console.</p>
Debug Trace	<p>Specifies whether SSL debug tracing should be enabled during the connection. Debug tracing provides more detailed messages than standard tracing.</p>
Verify Host Name	<p>Specifies whether you wish to verify that the host you are connecting to is the expected host. The host name in the host's digital certificate is compared against the value you supply in the Expected Host Name field. If the host name does not match the expected host name, the connection is refused.</p> <p>Note: The default context factories for TIBCO Enterprise Message Service automatically determine if host name verification is necessary. If you are using a custom implementation of the context factories, your custom implementation must explicitly set the verify host property to the correct value. For example:</p> <pre>com.tibco.tibjms.TibjmsSSL.setVerifyHost(false)</pre>

Field	Description (Cont'd)
Expected Host Name	<p>Specifies the name of the host you are expecting to connect to. This field is only relevant if the Verify Host Name field is also checked.</p> <p>If the name of the host in the host's digital certificate does not match the value specified in this field, the connection is refused.</p> <p>This prevents hosts from attempting to impersonate the host you are expecting to connect to.</p>
Strong Cipher Suites Only	<p>When checked, this field specifies that the minimum strength of the cipher suites used can be specified with the <code>bw.plugin.security.strongcipher.minstrength</code> custom engine property. See <i>TIBCO ActiveMatrix BusinessWorks Administration</i> for more information about this property. The default value of the property disables cipher suites with an effective key length below 128 bits.</p> <p>When this field is unchecked, only cipher suites with an effective key length of up to 128 bits can be used.</p>

JNDI Configuration

The JNDI Configuration shared resource provides a way to specify JNDI connection information that can be shared by other resources. This resource can be specified in any resource that permits JNDI connections. For example, [JDBC Connection on page 230](#) and [JMS Connection on page 235](#) can use JNDI connections. This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.</p> <p>This field appears on the New JNDI Connection Wizard. The name then appears in the title of the resource.</p>
Configuration		
Description	Yes	Short description of the resource.
JNDI Context Factory	Yes	<p>The initial context factory class for accessing JNDI. (<code>javax.naming.Context.INITIAL_CONTEXT_FACTORY</code>). You can choose from the drop down list of supported classes.</p> <p>Note: TIBCO BusinessEvents attempts to find the class. However, you may need to add the JAR file supplied by your JNDI service provider to the CLASSPATH environment variable to use JNDI.</p>

Field	Global Var?	Description (Cont'd)
JNDI Context URL	Yes	The URL to the JNDI service provider (<code>javax.naming.Context.PROVIDER_URL</code>). An example URL is provided when one of the supported JNDI context factory classes is selected. See your JNDI provider documentation for the syntax of the URL.
JNDI User Name	Yes	User name for logging into the JNDI server (<code>javax.naming.Context.SECURITY_PRINCIPAL</code>). If the JNDI provider does not require access control, this field can be empty.
JNDI Password	Yes	Password for logging into the JNDI server (<code>javax.naming.Context.SECURITY_CREDENTIALS</code>). If the JNDI provider does not require access control, this field can be empty.

Advanced Section

The Advanced section has the following fields.

Field	Global Var?	Description
Validate JNDI Security Context	No	Some application servers store the security context on the thread used to establish the JNDI connection (at the time of this release, only the WebLogic application server does this). In that case, the first activity to use this resource establishes the security context, then subsequent activities use the same security context, unless this field is checked. Checking this field ensures that each activity that uses this resource examines the security context to determine if the activity uses the same security context as the security context established on the thread. If they are different, the activity's configured security context is used. Checking this field causes additional overhead for activities that use this resource. Therefore, only check this field when necessary.

Field	Global Var?	Description (Cont'd)
Optional JNDI Properties	No	The table in this field contains optional properties to pass to the JNDI server. Use the +, X, and arrow keys to add, delete, and move properties in the list. Each property requires the property name, the datatype for the property, and the value for the property. See the documentation for your JNDI provider for more information about properties that can be passed to the JNDI server.

Rendezvous Transport

The Rendezvous Transport resource describes a TIBCO Rendezvous transport. See the TIBCO Rendezvous documentation for more information about specifying these fields.

This section provides a reference to the fields. For procedures see [Adding a Shared Resource on page 218](#).

Wizard and Configuration Tab

The Wizard and Configuration tab have the following fields.

Field	Global Var?	Description
Wizard		
Name	No	<p>The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323.</p> <p>Unlike other resource identifiers, however, shared resource identifiers can have spaces in the name.</p> <p>This field appears on the New Rendezvous Transport Wizard. The name then appears in the title of the resource.</p>
Configuration		
Description	Yes	Short description of the resource.
Daemon	Yes	<p>In the case of TIBCO Rendezvous daemon running on the same machine as TIBCO BusinessEvents engine, this is not specified. If Rendezvous is running on a different machine, then the Daemon field is specified as the remote host name followed by the socket number.</p> <p>For example:</p> <pre>ssl:acct:5785</pre>

Field	Global Var?	Description (Cont'd)
Network	Yes	This field contains the host name, IP address, network name, or interface name. For example: ; 224 . 34 . 103 . 4
Service	Yes	The TIBCO Rendezvous service name, specified in one of the following formats: <i>service name</i> or <i>port number</i> .
Use SSL	No	The Use SSL checkbox specifies that Secure Sockets Layer (SSL) should be used when communicating with the TIBCO Rendezvous daemon. When this field is checked, the Configure SSL button is enabled. See Configure SSL Button on page 246 for more information about configuring SSL parameters.

Configure SSL Button

The Configure SSL button allows you to configure the SSL parameters for communicating with the TIBCO Rendezvous daemon. See the TIBCO Rendezvous documentation for more information about how SSL is configured for TIBCO Rendezvous daemons and clients.

The SSL Configuration for TIBCO Rendezvous dialog has the following fields:

Field	Description
Daemon Certificate	File containing one or more certificates from trusted certificate authorities. This file is checked when connecting to a daemon to ensure that the connection is to a daemon that is trusted. This prevents connections to rogue TIBCO Rendezvous daemons that attempt to impersonate trusted daemons. You can retrieve a daemon's certificate using the administration interface in TIBCO Rendezvous. See the TIBCO Rendezvous documentation for more information about obtaining certificates through the administration interface. Once retrieved, you can select a folder in your project and import this certificate into the folder using the Tools > Trusted Certificates > Import Into PEM Format menu item.

Field	Description (Cont'd)
Identity	<p>This is an Identity resource used to authenticate to the TIBCO Rendezvous daemon. The Browse button allows you to select from a list of appropriately configured Identity resources. Only Identity resources whose Type field is set to Identity File or Username/Password are listed.</p> <p>See Identity Resource on page 228 for more information.</p>

Advanced Section

The Advanced section of the UI has the following fields.

Field	Global Var?	Description
RV Type	No	<p>The type of TIBCO Rendezvous connection to use. This can be reliable (standard RV transport), certified (RVCM), or Distributed Queue (RVCMQ).</p> <p>The fields of the Advanced tab correspond to the value selected for this field.</p>
Certified Transport		
CM Name	Yes	The name of the delivery-tracking session. This name is in the same format as TIBCO Rendezvous subject names.
Ledger File	Yes	The name and location of the persistent ledger file that tracks certified messages. If not specified, the certified message ledger is kept in process memory only.
Sync Ledger File	Yes	Specifies whether to keep the ledger file synchronous with the current messages.
Relay Agent	Yes	Name of the relay agent to use. Relay agents are useful when clients are disconnected from the network from time to time. The relay agents store inbound certified messages and labeled messages (and other messages related to certified delivery features) on behalf of their disconnected client programs. When a client is connected, it receives inbound messages immediately.

Field	Global Var?	Description (Cont'd)
Require Old Message	Yes	Check this box if you wish to require the retention of messages for which delivery has not been confirmed. These messages will be resent.
Message Timeout (sec)	Yes	The time limit (in seconds) for certified message delivery.
Distributed Queue		
CMQ Name	Yes	The name of the distributed queue. This name is in the same format as TIBCO Rendezvous subject names.
Worker Weight	Yes	The weight of the worker (this pertains to the worker processing queue requests). Relative worker weights assist the scheduler in assigning tasks. When the scheduler receives a task, it assigns the task to the available worker with the greatest worker weight.
Worker Tasks	Yes	Sets the task capacity for the worker (this pertains to the worker processing queue requests). Task capacity is the maximum number of tasks that a worker can accept. When the number of accepted tasks reaches this maximum, the worker cannot accept additional tasks until it completes one or more of them.
Worker Complete Time	Yes	The amount of time the scheduler waits for a worker process to complete. If the worker process does not complete in the specified period, the scheduler reassigns the message to another worker.
Scheduler Weight	Yes	Weight represents the ability of this member to fulfill the role of scheduler, relative to other members with the same name. Cooperating distributed queue transports use relative scheduler weight values to elect one transport as the scheduler; members with higher scheduler weight take precedence. Acceptable values range from 1 to 65535.
Scheduler Heartbeat	Yes	The scheduler sends heartbeat messages at this interval (in seconds). All members with the same name must specify the same value for this parameter. The value must be strictly positive.

Field	Global Var?	Description (Cont'd)
Scheduler Activation	Yes	When the heartbeat signal from the scheduler has been silent for this interval (in seconds), the member with the greatest scheduler weight takes its place as the new scheduler. All members with the same name must specify the same value for this parameter. The value must be positive

Chapter 16 **Rules and Rule Functions**

This chapter explains how to work with rules and rule functions.

For documentation on catalog functions see [Chapter 18, Functions, page 289](#)

Topics

- [Overview of Rules and Rule Functions, page 252](#)
- [Adding a Rule, page 254](#)
- [Rule Editor Reference, page 258](#)
- [Adding a Rule Function, page 260](#)
- [Rule Function Resource Reference, page 263](#)
- [Using Variables and Functions in the Rule Editor, page 265](#)
- [Using Priority and Rank to Control Order of Rule Execution, page 269](#)
- [Using the Quick Fix Feature in the Rule Editor, page 271](#)
- [Tips for Working in the Rule Editor, page 273](#)
- [Event Preprocessors, page 275](#)
- [Transaction Error Handler Rule Function, page 276](#)

Overview of Rules and Rule Functions

This chapter explains how to work with the rule editor to create rules and rule functions. It also explains some specific uses for rule functions.

To learn how rules are executed at runtime and for other background information see *TIBCO BusinessEvents Architect's Guide*, especially the section Understanding Conflict Resolution and Run to Completion Cycles.

See Also

- Chapter 4, Rules and Functions in *TIBCO BusinessEvents Architect's Guide*
- [Chapter 18, Functions, on page 289](#)
- [Chapter 19, Rule Language Grammar, on page 321](#)
- [Chapter 20, Rule Language Datatypes, on page 339](#)
- [Chapter 41, Diagrams, on page 701](#) for information on using rule dependency and sequence diagrams.

Form-based and Source Rule Editors

When you work with rules and rule functions, you can choose how to work:

- Using a form-based rule editor, similar to the rule editor in earlier versions of TIBCO BusinessEvents
- Using a source editor, which is closer to a Java programming environment.

You can switch between editors and changes made in one editor are reflected in the other one. You cannot switch from the source editor to the form editor if there are validation errors in the code.

Rule Components

A TIBCO BusinessEvents rule has three components:

- **Declaration** — Use the declaration to declare which concepts and events the rule will depend on, and the names by which instances of these entities can be referred to in the conditions and actions. Aliases must be valid identifiers. Declaring multiple terms of the same type allows the rule to consider multiple instances of the corresponding entity.

- **Conditions** — Each statement in the condition must evaluate to a boolean value. All of these statements must be true for the rule's action to be executed. Assignments and calls to certain functions are disallowed in the condition.
- **Actions** — List of statements that will be executed, when the rule is fired, for each combination of terms that matches all the conditions.

Effect of Cache Only Cache Mode

When using Cache Only cache mode for one or more entities, you must consider how to handle the cache-only entities when you write rules and preprocessor rule functions. See *Working With Cache Modes and Loading Cache Only Objects into the Rete Network* in *TIBCO BusinessEvents Architect's Guide*.

Adding a Rule

These instructions focus on the form-based editor and mention the equivalent settings in the source editor. Adapt the instructions if you are using the source editor.

Figure 2 Rule Form Editor

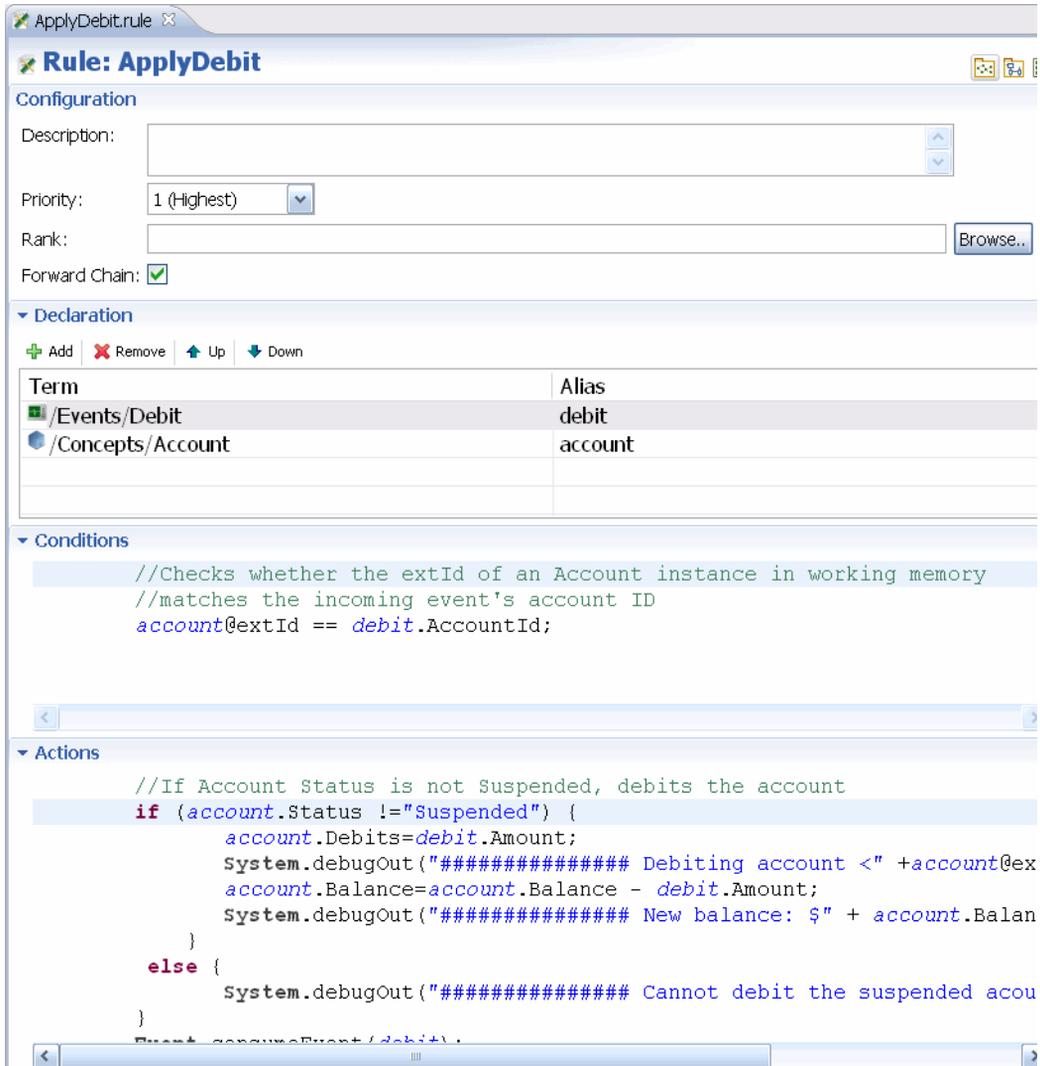


Figure 3 Rule Source Editor

```

* @description
* @author
*/
rule Rules.ProcessDebits.ApplyDebit {
    attribute {
        priority = 1;
        forwardChain = true;
    }
    declare {
        Events.Debit debit;
        Concepts.Account account;
    }
    when {
        //Checks whether the extId of an Account instance in working memory
        //matches the incoming event's account ID
        account@extId == debit.AccountId;
    }
    then {
        //If Account Status is not Suspended, debits the account
        if (account.Status != "Suspended") {
            account.Debits=debit.Amount;
            System.debugOut("##### Debiting account <" +account@extId+ "> by $"
+debit.Amount);
            account.Balance=account.Balance - debit.Amount;
            System.debugOut("##### New balance: $" + account.Balance);
        }
        else {
            System.debugOut("##### Cannot debit the suspended account <"
+account@extId +">");
        }
        Event.consumeEvent(debit);
    }
}

```

Optionally add entry for Rank as needed (Rank=RuleFunction), or enter in Form view.

To Add a Rule

See [Rule Editor Reference on page 258](#) for details on the settings.

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the rule and select **New > Rule**. You see the New Rule Wizard.

2. In the Rule Name field, type a name for the rule. In the Description field, type a description as desired. (In the source editor the description appears in the `@description` line of the comments at the top of the editor.)



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. Click **Finish**. By default, you see the source rule editor on opening the editor. It shows the outline for a rule's source code. Click the **Form** tab at the bottom of the editor to use the form editor.

At any time you can click the Form and Source tabs at the bottom of the editor to switch between the form-based editor and the source editor as you work, depending on your preference.



Rule Editor Preference To set which mode the editor uses upon first opening, go to `Window > Preferences > TIBCO BusinessEvents > Rules` and check or uncheck the Initially show 'Form' tab checkbox as desired.

4. In the Form editor Configuration section, add or edit a description as desired. (In the source editor the description appears in the `* @description` line of the comments at the top of the editor.)
5. To control the order of rule execution:
 - a. If you want to control the order in which rules execute, set the Priority field accordingly. Highest priority is 1. (In the source editor set `priority = n` where `n` is the priority number.)

If you want to also control the order in which rules with the same *priority* execute, set the Rank field accordingly. Browse to and select the rule function you created for this purpose. (If you want to use the source editor, add `Rank=RuleFunction` in the list of rule attributes.) See [Using Priority and Rank to Control Order of Rule Execution on page 269](#) for details.
6. To disable forward rule chaining (the default behavior) uncheck the Forward Chain checkbox. (In the source editor set `forwardChain = false`.)

7. In the Declarations section (equivalent to the `declare` statements in the source editor), drag an ontology entity into the section, OR do the following:
 - a. Click **Add** to add resources that you will be using in your rule. You see the Select Rule Declaration Arguments dialog.
 - b. In the upper half of the Select Rule Declaration Arguments dialog, select the kind of entity you want to use.
 - c. In the lower half of the dialog, select a resource from the filtered ontology tree, and click **OK**. Your selection appears in the Declarations list. TIBCO BusinessEvents assigns an alias to the resource. You can edit the alias.
 - d. To re-order the declarations, highlight a declaration and click the up or down arrow to move it. This is relevant only in rule functions, to order the arguments.

Repeat to add more entity types as desired.

8. In the Conditions section (equivalent to the `when` statements in the source editor), write condition statements (in the TIBCO BusinessEvents rule language).

Each line is a complete statement. Each condition must evaluate to a Boolean value. Each line is joined to the others with an implicit AND operator. All of a rule's conditions must evaluate to true for the conditions to be satisfied. See *Order of Evaluation of Rule Conditions* in *TIBCO BusinessEvents Architect's Guide* for more information.

See [Using Variables and Functions in the Rule Editor on page 265](#) and [Tips for Working in the Rule Editor on page 273](#) for more information on working in the rule editor.

9. In the Actions section (equivalent to the `then` statements in the source editor), write action statements (in the TIBCO BusinessEvents rules language).
10. Save the resource.

Rule Editor Reference



The rule editor and rule function editor are similar. This section focuses on the form-based rule editor. You can adapt the information in the Area and Property section to apply to the different blocks of code in the source editor.

Property	Description
Configuration Section	
Name (Wizard only)	The name to appear as the label for the resource. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
Description (Editor and Wizard)	Short description of the resource.
Priority (Editor and Wizard)	Specify a value between 1 and 10, where 1 is the highest priority and 10 is the lowest priority. Rules with a higher priority execute before rules with a lower priority. See also Rank. Only set priority or rank where there is a reason to control the order of rules. As a general practice, it is more efficient to let the engine determine the order of rule firing.
Rank	Specifies a rule function that controls the order of execution of rules with the same priority. Validity must include Condition (that is, do not select Action Only as the validity.) See Using Priority and Rank to Control Order of Rule Execution on page 269 for details. Default is 0.0.
Forward Chain	Determines if the rule is used in forward chaining. If the checkbox is unchecked, changes made by this rule won't trigger other rules. Default is checked.

Property	Description
Declaration Section	
Term	<p>A concept or event type in the project that you will use in your rule. Types you add to the declaration define the scope of the rule.</p> <p>For example, a concept definition such as <code>/Concepts/Accounts/CheckAccount</code>.</p> <p>It is not necessary to add scorecards in the declaration in order to use them in the rule.</p>
Alias	<p>A name used to refer to the scope element in the body of the rule. You can change the alias. Like resource names, aliases must be valid identifiers See Identifier Naming Requirements on page 323.</p>
Conditions Section	
	<p>Each line in the Conditions area is a single expression that evaluates to <code>true</code> or <code>false</code>. Each line is joined to the others with an implicit and operator.</p> <p>For the OR operator, use a double pipe (<code> </code>) on the same line.</p> <p>TIBCO BusinessEvents evaluates single conditions from left to right. TIBCO BusinessEvents optimizes the evaluation of multiple conditions (see Order of Evaluation of Rule Conditions in <i>TIBCO BusinessEvents Architect's Guide</i>.)</p>
Actions Section	
	<p>List of statements that will be executed when the rule is fired.</p>

Adding a Rule Function

Regular rule functions have arguments and a body containing the code for the function. Virtual rule functions have arguments but no body. Virtual rule functions are used only with the TIBCO BusinessEvents Decision Manager add-on software. The implementation for virtual rule functions is provided by one or more decision tables. See TIBCO BusinessEvents Decision Manager product documentation for more details.

To Add a Rule Function

See [Rule Function Resource Reference on page 263](#) for details on completing values.

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the rule function and select **New > Rule Function**. You see the New Rule Function Wizard.
 - a. In the Rule Function Name field, type a name for the rule function.
 - b. In the Description field, type a description. (In the source editor the description appears in the * @description line of the comments at the top of the editor).
 - c. Set the return type for the rule function. Default is `void`. Browse to select a return type as needed. For return types that require additional configuration, such as `ContainedConcept`, complete the configuration in the Rule Function editor.
 - d. If you want this rule function to be a virtual rule function (to be implemented by a decision table), check the **Virtual** checkbox.



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

2. Click **Finish**. If you see the source editor, click the **Form** tab at the bottom of the editor to use the form editor as desired.

At any time you can click the Form and Source tabs at the bottom of the editor to switch between the form-based editor and the source editor as you work, depending on your preference.



Rule Function Editor Preference To set which mode the editor uses upon first opening, go to Window > Preferences > TIBCO BusinessEvents > Rules and check or uncheck the following checkbox as desired: **Initially show 'Form' tab in Rule Function Editor.**

3. In the Form editor Configuration section, add or edit an alias and a description as desired. (In the source editor the description appears in the * @description line of the comments at the top of the editor and the Alias appears in the attribute list.)
4. If you did not do so in the Wizard, use the Return Type to select the return type of the rule function.
5. If you did not do so in the Wizard, set the Virtual check box according to your need. Check the check box if you are creating a virtual rule function (to be implemented by a TIBCO BusinessEvents Decision Manager decision table).



In the source editor, the signature of a virtual rule function is:

```
virtual void rulefunction folder.RFName
```

Do not add code to the Body block in the source editor of a virtual rule function. If you do, you see error messages if you try to save or to switch to the form-based editor.

6. From the Validity drop-down list, select the value that specifies where the rule function can be used (source editor attribute equivalents shown in parentheses):
 - Action (validity=ACTION)
 - Action and Condition (validity=CONDITION)
 - Action, Condition and Query (validity=QUERY)

Virtual rule functions have a non-editable validity setting of Action.

7. If the rule function returns a value, specify the Return Type, otherwise leave this field set to void. (Appears in the signature of the rule function in the Source editor.)

Virtual rule functions have a non-editable return type of Void.

8. In the Scope section (scope statements in the source editor) you define the arguments of the rule function. Drag entities into the Scope area from BusinessEvents Studio Explorer, OR do the following:
 - a. Click **Add** to add resources that you will be using in your rule function. You see the Select Rule Function Scope Arguments dialog.
 - b. In the upper half of the Select Rule Function Scope Arguments dialog, select the type you want to use.
 - c. If you want to specify an array, check the **isArray** checkbox. (You can specify a variable array in the source editor in the usual way, for example, `int[] myArr`.)
 - d. If the type you select is an ontology type, in the lower half of the dialog, select a resource from the filtered ontology tree.
 - e. Click **OK**.

Your selection appears in the list. TIBCO BusinessEvents assigns an alias to it. You can edit the alias.

Add more entities as needed.

9. Add more arguments as needed, and use the up and down arrows to order the arguments as needed.
10. In the Body section (Body statements in the source editor), use the TIBCO BusinessEvents rule language to implement the function. (Virtual rule functions have only a signature, and no implementation at design time.)
See [Using Variables and Functions in the Rule Editor on page 265](#) and [Tips for Working in the Rule Editor on page 273](#) for more information on working in the rule editor.
11. Save the resource.

Rule Function Resource Reference



Rule Function resources enable you to write rule functions that you can use in rules, as startup and shutdown actions, and as preprocessors.

Virtual rule functions are decorated with a V.

Property	Description
Configuration Section	
Name (Wizard only)	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
Description (Editor and Wizard)	Short description of the resource.
Return Type (Editor and Wizard)	If the rule function returns a value, specify the Return Type, otherwise leave set to void. See the list of valid types for the Scope Section area, next.
Virtual (Editor and Wizard)	If set to yes, the rule function is a virtual rule function. Virtual rule functions have arguments but no body. The Body section is disabled and so is the Return Type field. The body is provided by a decision table. See <i>TIBCO BusinessEvents Decision Manager User's Guide</i> .
Alias	Optionally, enter an alias for the rule function. Used as a short way to refer to the rule function. You can use the alias, for example, to make query strings shorter (if you have the TIBCO BusinessEvents Event Stream Processing add-on).

Property	Description
Validity	<p>Specifies where the rule function can be used. Possible values are as follows:</p> <p>Action Indicates that this rule function can be used only in the Action block of a rule.</p> <p>Action and Condition Indicates that this rule function can be used in the Action and Condition blocks of a rule.</p> <p>Action, Condition and Query Indicates that this rule function can be used in the Action and Condition blocks of a rule, and can also be used in the text of a query (The query language features are available only in TIBCO BusinessEvents Enterprise Suite).</p> <p>Note: Only Action rule functions can be used as startup rule functions or shutdown rule functions.</p>

Scope Section

Term	<p>The type of the argument. Arguments and return type can be any of the following, including arrays of these datatypes:</p> <ul style="list-style-type: none"> • Primitive, that is any of: <code>String</code>, <code>int</code>, <code>long</code>, <code>double</code>, <code>boolean</code>, <code>DateTime</code>, <code>Object</code> • Concept • Event • Specific type of Concept • Specific type of Event <p>The Object data type is used to pass parameters between standard and user-defined functions and external Java sources.</p> <p>For more details, see Chapter 20, Rule Language Datatypes, on page 339.</p>
Alias	<p>Each argument requires a type and an alias. Names must be valid identifiers. See Identifier Naming Requirements on page 323.</p>

Body Section

List of statements that will be executed when the rule function executes.

Using Variables and Functions in the Rule Editor

This section provides some tips on working in the rule editor. The rule editor is used for TIBCO BusinessEvents rules, rule functions, and state machine rules.

Using Catalog Functions in the Rule Editor

To use catalog functions in an editor choose one of the following methods.

1. When adding code in the rule editor, do one of the following.
 - Begin to type the function name, beginning with the folder path to the function.

Do not enter the catalog name (which appears in the list of catalog functions in the Catalog Functions view, and looks like a top-level folder). Type a period at the end of the folder name. A popup window shows all folders and functions within the folder whose name you typed. For example type `Database.` to see a list of all functions in the Database folder.
 - Open the Catalog Functions view. To open the view click as follows:

Window > Show View > Other > TIBCO TIBCO BusinessEvents > Catalog Functions.

Drill down on categories within the catalog to expand to lists of functions and drag the desired function to the rule editor.
2. Provide parameter values as indicated by the tooltip. You can hover the cursor over a function to display a tooltip showing the function's arguments. You can also see the tooltip contents in the online reference, *TIBCO BusinessEvents Functions Reference*.

See [Overview of Catalog Functions on page 290](#) for descriptions of the various function catalogs, and an explanation of the decorations that appear on many function names.

For information about configuring mapper functions see [Using the Function Argument Mapper on page 266](#).

Using Global Variables in the Rule Editor

To use a global variable in the rule editor, use one of the `System.getGlobalVariableAs*` functions. For example:

```
System.getGlobalVariableAsString("Hostname", "Localhost")
```

Where `Hostname` is the name of the variable and `Localhost` is an optional literal

value to use if the variable is not found.

Do not use this syntax: `%%Global.Variable.Name%%`.

See [Working with Global Variables on page 14](#) for more details about global variables.

Using the Function Argument Mapper

For functions known as mapper functions you can use the Function Argument Mapper to map inputs from a source to the function arguments.

See [Chapter 21, Mapping and Transforming Data, on page 345](#) for a reference to using the Function Argument Mapper.

To Open the Function Argument Mapping Wizard

You can open the Function Argument Mapping Wizard in various ways. For example here is one way:

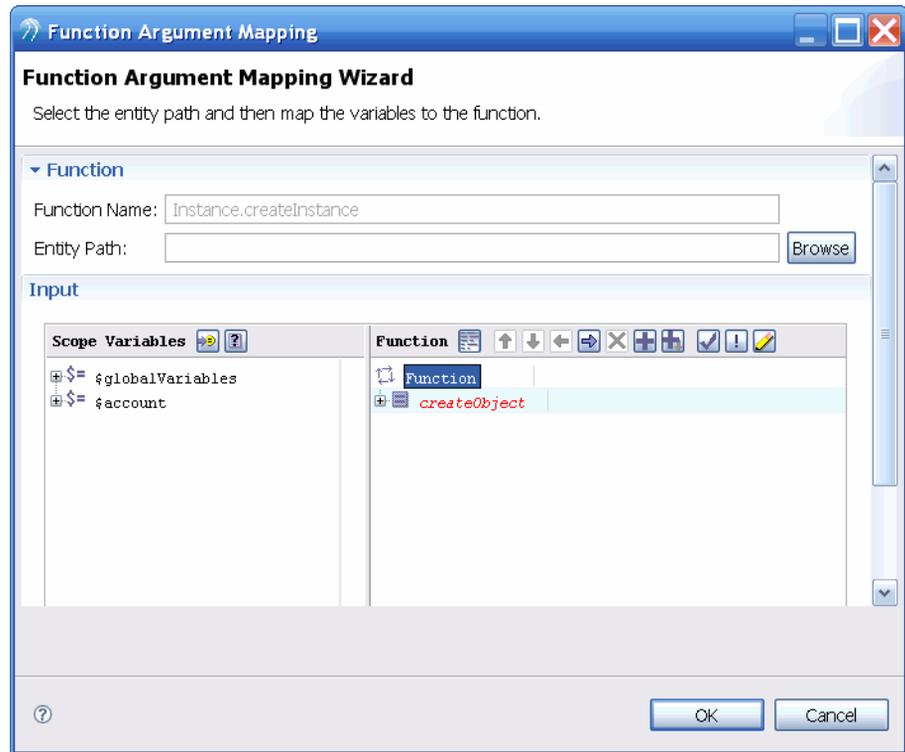
1. In the rule editor, type the category of function you want to use. A list of functions in that category appears.
2. Select a mapper function from the list of functions in that category. The function argument area contains the text `"xslt://"`.
3. Control-click the text `"xslt://"` to display a hypertext link. Click the link to open the Function Argument Mapping Wizard.

Other ways to perform the same action are as follows:

- Type the name of a mapper function category into the **Actions** or **Conditions** areas, then type an open parenthesis ("(") TIBCO BusinessEvents displays `"xslt://"`. Control-click to open the Function Argument Mapping Wizard.
- Type the entire string to specify the function category path and name, followed by (`"xslt://"`). Then control-click the text `"xslt://"` to open the Function Argument Mapping Wizard.

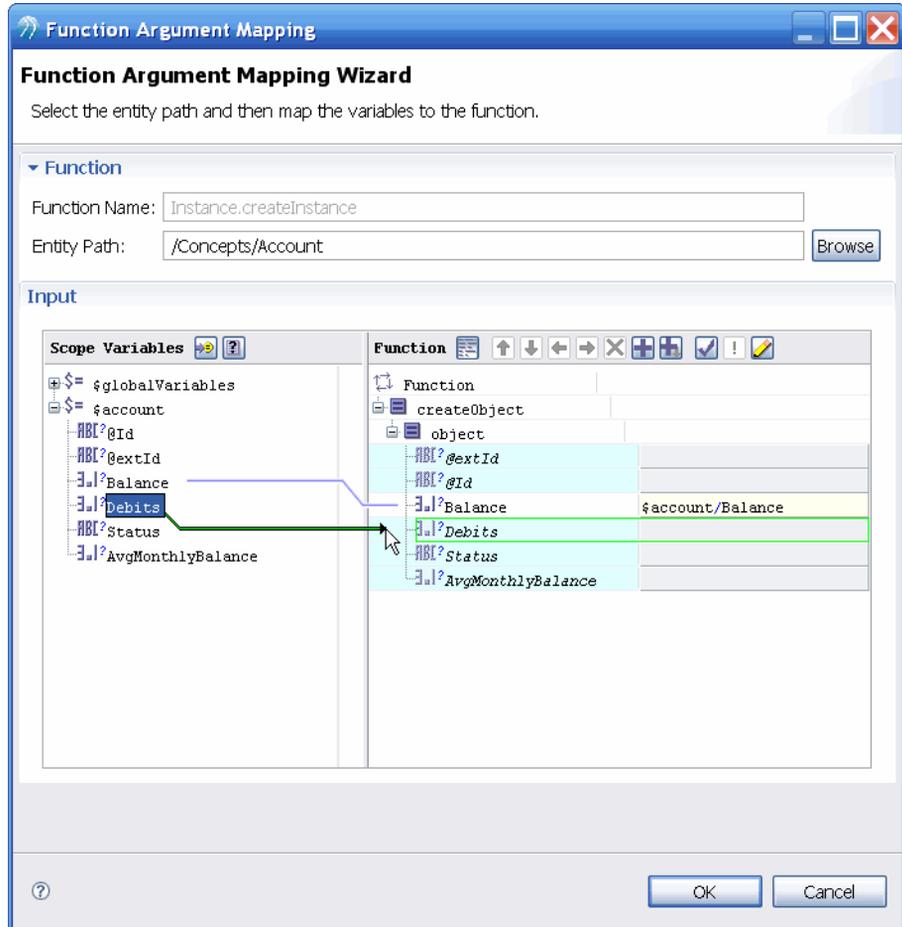
To Work with the Function Argument Mapper Wizard

See [Chapter 21, Mapping and Transforming Data, on page 345](#) for full details on using the wizard. This section provides summary information only.



1. In the Function section, take appropriate action. For example, if you are using `Instance.createInstance()`, you would click the browse icon () to the right of the **Entity Path** field and select a concept type so you can create an instance of it.
2. Select a resource and click **Apply**.
3. Select the **Input** tab. TIBCO BusinessEvents displays a list of variables associated with the project on the left and a list of properties for the selected resource type on the right.
4. Drill down and expand the lists on both sides to expose the variables and properties.

5. Drag variables from the left to the appropriate property on the right, as shown:



6. If you want to define the arguments using more complex logic, type the code or click the XPath Formula Builder () button to use the XPath Formula Builder. See [Chapter 22, XPath Formula Builder, on page 383](#).
7. Click **OK**.



Array index difference In the TIBCO BusinessEvents language, array indexes start from zero (0). However, in XSLT and XPath languages, they start from one (1). It's important to remember this difference when using the rule language in the rule editor, and when working in the XSLT mapper and the XPath builder.

Using Priority and Rank to Control Order of Rule Execution

For each RTC, the rule agenda is sorted by priority and then within priority by rank, for those rules that use the same ranking. Use of priority and rank is optional. You can also use priority without using rank.

In the rule's Rank field (or rule attribute, in the source view), you specify a rule function that returns a double. The larger the return value, the higher the ranking. You can specify the same rule function in different rules to perform ranking across tuples of those rules. Here are the requirements:

- The rule function must have a Validity setting that includes Condition (that is, do not set it to Action Only).
- You can assign the same rule function to different rules as long as the following is true:
 - The scope of the rule function includes only parameters found in all the rules that use the same function. This can be a subset of what's in the rule declaration. (As with rules, primitives are not allowed.)
 - The parameters must be used in the same order as they appear in the rule declaration.
- The rule function must return a double value. (The default value for the Rank field is 0.0.)

Examples

For example, suppose two rules declare CustName and SupportLevel (plus other things that can be different). You assign the same rule function to both rules. The function returns 3.0 for Gold support level, 2.0 for Silver, and 1.0 for Bronze. As a result, among rule tuples with priority 1, those for customers with Gold support execute before those for customers with Silver support, which execute before those for customers with Bronze support.

Below is an example showing how rule priority and ranking determine the sort order. Suppose you assign the same rule function for ranking rules 1, 2, and 3. At run-time these are some rule tuples to be sorted in the agenda for an RTC:

Rule 1 (Tuple X)	Priority: 5	Rank: 10.0
Rule 2 (Tuple A, B)	Priority: 1	Rank: 1.0
Rule 2 (Tuple A, C)	Priority: 1	Rank: -1.0
Rule 2 (Tuple A, D)	Priority: 1	Rank: -2.0
Rule 3 (Tuple A)	Priority: 4	Rank: 0.0

They are sorted and executed as follows. (This could change during a conflict resolution cycle, of course, depending on the effect of rule actions.)

Rule 2 (Tuple A, B)	Priority: 1	Rank: 1.0
Rule 2 (Tuple A, C)	Priority: 1	Rank: -1.0
Rule 2 (Tuple A, D)	Priority: 1	Rank: -2.0
Rule 3 (Tuple A)	Priority: 4	Rank: 0.0
Rule 1 (Tuple X)	Priority: 5	Rank: 10.0

Reverse Order Example

If you want the rule with the smallest return value to be ranked highest, multiply by negative one (-1.0) to reverse the size of the values. For example suppose you want rules that were asserted earlier to execute before those that were asserted later. The time values returned by three rules are 100, 150, and 300. Using $-1.0 * time$, they will be fired in the desired order: -100.0 (the largest value), then -150.0, then -300.0.

Using the Quick Fix Feature in the Rule Editor

The Quick Fix feature enables you to create concepts, events, and rule functions without leaving the rule editor. It also enables you to add properties to existing concepts and events. The feature is available when an unknown reference appears in the rule or rule function code.



This feature is also used in the state modeler feature, available in TIBCO BusinessEvents Data Modeling, wherever the rule editor is used.

If a qualified name has multiple unresolved references, for example, `Concepts.Something.SomethingElse`, the Quick Fix feature only applies to the first unresolved reference in the name.

To Use the Quick Fix Feature

1. Type the code in the rule editor as if the entity, property, or rule function you want to use already exists. See [Table 28, Quick Fix Feature Options, on page 271](#) for some examples.

You see a lightbulb icon (💡) in the left margin on any line where there are unknown references that you can configure using Quick Fix.

2. To use the available Quick Fix options, click the lightbulb icon or press Ctrl+1. The options are explained in [Table 28](#):

Table 28 Quick Fix Feature Options

Unknown Reference	Quick Fix Editor
Unqualified name. Example: <code>SomeName s;</code>	A list of appropriate wizards appears so you can create a new concept, simple event, or time event. The wizard you select opens so you can define the entity.
Unqualified name followed by parentheses. Example: <code>SomeName()</code>	A link to the rule function creation wizard appears. Click it to open the wizard and define the rule function.
Qualified name. Examples: <code>Rules.SomeName();</code> <code>Concepts.SomeName s;</code>	A Quick Fix option appears so you can automatically create the entity or rule function in the given folder. The folder must already exist in the project.

Table 28 Quick Fix Feature Options

Unknown Reference	Quick Fix Editor
Qualifier is a concept or event and reference is unknown. Example:	A Quick Fix option appears with the name: Create a property definition in the entity ' <i>entityName</i> '. Click the option to open the New Property Definition dialog, and configure the details of the property.
<code>Concepts.Person person; person.someProperty = "123"</code>	In the example, the <code>someProperty</code> property does not already exist as a property of <code>Concepts.Person</code> .

Tips for Working in the Rule Editor

This section has some tips to help you work in the rule editor.

Switching between Form and Source Editors

You can freely switch between form and source editors for rules and rule functions. In each case the editors stay synchronized with the latest changes.

You cannot switch from the source editor to the form editor if there are any syntax or resolution errors in your code.

The Priority setting is used by the runtime engine when determining the order in which rules are fired. Rules with a number closer to one fire first. When there is no reason to force rules to execute in a particular order, leave the Priority set to the default and let the runtime engine determine rule order.

Declaring multiple terms of the same type Allows the rule to consider multiple instances of the corresponding entity. Specify different aliases to keep the terms distinct

Scorecards Scorecards are like concepts except that there is only one instance of a scorecard (or more accurately, one instance per agent when multi-engine features are used). It is therefore not necessary to put scorecards in the declaration of a rule because a scorecard never requires an alias. You can use scorecard properties in conditions (just as you would concept properties). However, because a scorecard doesn't have an alias, refer to it like a function, for example,
`Folder.Folder.Scorecard.prop1`

Standard Eclipse Features

In addition to some TIBCO BusinessEvents-specific features, the source and form rule editors support standard Eclipse functionality such as the following: Undo and redo; copy and paste; breakpoint features; standard text annotations (which can be changed using Preferences); text folding (source editor only); Java outline view (source editor only).



When you're working in the source editor, press **Ctrl+Shift+L** to see a list of keyboard shortcuts available in that context. (This is a general Eclipse feature.)

Information Highlighted

Keywords, variables, and functions are highlighted in the text. Also, when you hover the mouse over a resource such as a concept, event or function, information about it displays in a tooltip.

Syntax and resolution errors are automatically flagged by visual cues. They are underlined and also display in the vertical and overview rulers.

Miscellaneous Tips

Some other tips are highlighted below.

Table 29 Tips for Working in the Rule Editor

To do this...	Do this...
Switch Between Source and Form Editors	Click the bottom tabs to switch between the source and form editors. The code remains synchronized. However, you cannot switch if there are errors in the code. First resolve the errors, then switch.
To Use Content Assist to Complete Values	<p>The content assist feature helps you complete values using information that is available in resources. For example, if you type the name of a concept type (or its alias) and then a period, a list of the concept type properties appears for you to select from.</p> <p>The selection list also appears when the cursor is in an appropriate location and you press Ctrl+Space, or if you right click and select Edit > Content Assist from the context (right-click) menu.</p>
To Comment (and Uncomment) a Line	To comment out a line, or uncomment a line, press Ctrl+/ or select Edit > Toggle Comment from the context menu
To Search for References	When the cursor is placed in an appropriate item in the code such as an entity or function name, you can find all references to that item references in the rule code. Press Ctrl+Shift+G , or select Search > Search for References from the context (right-click) menu. The item references are highlighted in the text, and an arrow appears in the vertical ruler.
To Jump to the Definition of an Item	<p>To jump to the location where an item is defined, you can use two methods.</p> <ul style="list-style-type: none"> Click in the item name and press F3, or right-click and select Open Declaration. Press and hold the Control key while you move (hover) the mouse pointer the text. When you hover over an item that displays an underline, Ctrl+click the item to jump to the place where it is defined. <p>For example, you would jump from an alias to the declaration, and from an entity or entity property to the entity's editor.</p>

Event Preprocessors

Event preprocessors are rule functions with one argument of type simple event. They perform tasks after an incoming message is transformed into a simple event but before it is asserted into the Rete network.

An event preprocessor is assigned to a destination and acts on all events arriving at that destination.



See Event Preprocessors in *TIBCO BusinessEvents Architect's Guide* for important information on the use of preprocessors.

If an event preprocessor fails due to an exception, the Retry on Exception field setting on the event type determines subsequent behavior (see [Simple Event Reference on page 167](#)).

Loading Cache Only Entities If you are using the Cache Only cache mode for any entities, it is recommended that you load the relevant entities in the event preprocessor. See [Cache Related Functions on page 309](#).

Configuring an Event Preprocessor

To configure a preprocessor you associate the desired rule function with a destination using the Cluster Definition Descriptor editor. See [Configuring Collections of Rules, Rule Functions, and Destinations, page 492](#).

Transaction Error Handler Rule Function

You can create a transaction error handler callback rule function that enables you to identify which post RTC transactions failed or which events were not sent out during the post RTC phase. The transaction error handler rule function is invoked each time a database transaction exception occurs, and each time a send event exception occurs.

This feature requires cache-aside database write strategy. When backing store is enabled, the default value for the property `Agent.agentclassname.enableParallelOps` is true. The feature is not supported when the property value is set to false:

```
Agent.agentclassname.enableParallelOps=false
```



When the backing store is disabled, the default values for the cache-aside database write strategy and the property `Agent.agentclassname.enableParallelOps` are false.

Set these two values to true explicitly in the CDD file in order to use the transaction error handler rule function.

The error handler rule function provides full details of the failed transactions. The transaction details provided in this callback function can be used for audit trail purposes. The engine continues retrying the transaction or action as usual.

To Register the Rule Function

To register the rule function, add the following property to the `be-engine.tra / project` CDD file at an appropriate level (for example, at the cluster level to affect all deployed engines). Set the value to the project path of the error handler rule function:

```
be.engine.txn.error.function=projectPathToErrorHandlerRuleFunction
```

Rule Function Signature

The required signature for the error handler rule function is:

```
int Name (Object txns, int errorType, int errCode, String errMsg, long retryCount)
```

The return value must be 0 (zero).

The following table explains the arguments.

Argument	Notes
Object <i>txns</i>	<p>You can turn this argument into an <code>Object[]</code> within the rule function, for example:</p> <pre>Object[] array = txns;</pre> <p>The resulting array can be used to obtain useful data:</p> <ul style="list-style-type: none"> • <code>array[0]</code> is a <code>Concept[]</code> containing all the <code>Concept</code> objects that: <ul style="list-style-type: none"> — have been created in the current RTC, — and have not been deleted. • <code>array[1]</code> is a <code>Concept[]</code> containing all the <code>Concept</code> objects that: <ul style="list-style-type: none"> — existed before the current RTC, — have been modified in the current RTC, — and have not been deleted. • <code>array[2]</code> is a <code>Long[]</code> containing the id of all the <code>Concept</code> objects that: <ul style="list-style-type: none"> — existed before the current RTC, — and have been deleted in the current RTC. • <code>array[3]</code> is a <code>SimpleEvent[]</code> containing all the <code>SimpleEvent</code> objects that: <ul style="list-style-type: none"> — are going to be added to the cache. • <code>array[4]</code> is a <code>Long[]</code> containing the id of all the <code>Event</code> objects that: <ul style="list-style-type: none"> — existed before the current RTC, or arrived and started the RTC, — and have been deleted in the current RTC.
int <i>errorType</i>	<p>Value can be one of:</p> <ul style="list-style-type: none"> -1 for errors happening during database operations. -2 for errors happening when sending an event.

Argument	Notes
<code>int <i>errCode</i></code>	<p>The error code, which is dependent on <code>errorType</code>:</p> <ul style="list-style-type: none">• For database related errors:<ul style="list-style-type: none">— The error code as returned by <code>SQLException</code>, if available— Else -100 if TIBCO BusinessEvents determines that the database is not available— Else -200, which means that this is an unknown error• For sent event errors, the value is -1, meaning an internal error condition.
<code>String <i>errMsg</i></code>	<p>The associated exception message.</p>
<code>long <i>retryCount</i></code>	<p>The number of times this transaction has been retried before the present call to the callback.</p>

Chapter 17 **Rule Templates**

Rule templates and rule template views are used for creating business rule in TIBCO BusinessEvents WebStudio. This chapter explains what they are for and how to create them. See *TIBCO BusinessEvents WebStudio User's Guide* for more information on business rules.

Topics

- [Rule Template and Rule Template View Overview, page 280](#)
- [Adding a Rule Template, page 284](#)
- [Adding a Rule Template View, page 286](#)
- [Rule Template Editor Reference, page 287](#)

Rule Template and Rule Template View Overview

A rule template is a specialized type of rule, and a rule template view puts a user-friendly interface around the rule template for use in Web Studio.

Rule templates and rule template views enable non-technical Web Studio users to define executable rules (called *business rules*) within limits defined in the rule template.

This chapter assumes familiarity with rules. For details on creating rules, see [Chapter 16, Rules and Rule Functions, on page 251](#)

The Rule Template Editor

Comparing similar sections of the Rule Template editor and the Rule editor highlights the unique features of rule templates.

Configuration Section

The Configuration section is the same as the Configuration section in a rule. Priority, ranking, and forward chaining features work the same at runtime whether a rule is created using the Rule Editor or the Rule Template Editor.

Declarations and Variables Section

The Declarations and Variables section (`declare` in the source view) is similar to the Declaration section of a rule. However, in rule templates you can declare primitive types (as you can in rule functions). You can also define initial expressions for primitive types, for example, `int j = 50;`

Pre-conditions Section

The Pre-conditions section (when in the source view) is similar to the Conditions section of a rule. However, WebStudio users can add additional conditions when defining individual business rules. The Pre-conditions section specifies conditions that must be met in *all* instances of a rule template before the rule's actions execute.

Another difference is that the Pre-conditions section can contain non-conditional statements, for example, local variable declarations, to provide more flexible and comprehensive condition checking capabilities. Here is a simple example of such a section:

```
int i=con.prop1;  
int j=con2.prop1;
```

```
i > j;
i == j;
```

This feature enables complex calculations to be done before the definition of the condition. The Pre-conditions section could also contain conditions that might appear in any rule, such as the following:

```
con.prop1 > con2.prop1;
```

Action Context Section

The Action Context section (`actionContext` in the source view) has a similar purpose to the Actions section of a rule. The Action Context section, however, defines *all possible actions* that can be taken by a business rule (after all conditions are met). Only the action context statements that the WebStudio user selects and defines as commands in the business rules are actually taken (depending on rule evaluation at runtime).

Defining the superset of all possible actions simplifies and restricts the configuration of business rules in WebStudio.

The actions are not completely defined in the rule template. Completing the definition of an action is a WebStudio user task. If bindings are used (and a view) then in WebStudio, the business rule writer has to enter only the binding values to complete the definition.

Action context statements are of three types: create, modify, and call, plus arbitrary actions, as explained next.

Create Statements

Create statements enable business rule developers to create new concepts or events. Create statement uses fully qualified name for the concepts and events instead of simple name. For example:

```
create Concepts.Bulk bulk;
```

In the rule template, that is as far as you go. It's up to the business rule developer to assign the specific constructor values for the entity type.

Modify Statements

Modify statements enable business rule developers to modify concepts or events. For example:

```
modify bulk;
```

Only entities in the scope of the rule can be used. As with create statements, it is up to the business rule developer to determine which properties to change and assign values to those properties.

Call Statements

Call statements enable business rule developers to call rule functions. Call statement uses fully qualified name for the concepts and events instead of simple name. For example:

```
call RuleFunctions.matchFound matchfound;
```

The business rule developer assigns the parameters for the rule function.

Arbitrary Actions

Rule template developers can add arbitrary actions as well as create, modify, and call statements. For example:

```
create Concepts.Bulk bulk;
if (mybinding > 10) {
    MyRFs.myRF(otherbinding);
    mycon.prop1 = thirdbinding;
}
```

Bindings Tab

The form editor for rule templates provides an additional tab, the Bindings tab, with two sections that are not found in the Rule editor: Bindings and Views (bindings and views in the source view).

Bindings

A binding is used like a local variable whose value is assigned by WebStudio users when defining business rules.

When defining a rule template, you can use bindings in condition checks and in action statements. For example, suppose you define a binding called `dollarAmount`. You could then define a condition as follows:

```
output.UNIT_PRICE > dollarAmount;
```

At runtime, the `UNIT_PRICE` is checked against the specific dollar amount defined in the business rule.

Similarly, you could use the binding in an action statement as follows:

```
RuleFunctions.matchFound(match, null, dollarAmount);
```

When executed, the engine calls the `matchFound` rule function and passes the value of `dollarAmount` as defined in the business rule.

Bindings are primitive types. You can optionally provide an initial value, and the value can be defined using a domain model. If a domain model is specified for a binding, the view defined for a rule template displays the input field for the binding as a drop down box. In the following example, the String binding `day` is initially assigned to the value of `Monday`, and it is tied to the domain model `RuleFunctions.WeekDays`.

```
String day = "Monday" (RuleFunctions.WeekDays);
```

Views In the Views section of the **Bindings** tab, you select a view that you have already defined. The selected view is used in WebStudio to present the rule template to the users and enable them to define the business rules (that is, executable rules).

If you do not associate a view with a rule template, a builder-based interface is used.

Rule Template Views

A rule template view defines a visual presentation of the rule template to make it easy for the WebStudio user to define a business rule.

The View uses rule template bindings to define an easy-to-use form-like interface where the WebStudio user assigns values to the bindings, for example, threshold values. For example:

```
Require that all applicants have a minimum income of <binding
id="minimumIncome"></binding>, a minimum age of <binding
id="minimumAge"></binding>, and restrict the credit type to
<binding id="creditType"></binding>
```

If the binding is associated with a domain model, the WebStudio interface will display the input widget as a drop down list. Otherwise, an input widget of the appropriate type is displayed, for example, a check-box for a boolean field.

Adding a Rule Template

See [Rule Template and Rule Template View Overview on page 280](#) for an introduction to this feature, and [Rule Template Editor Reference on page 287](#) for a reference to the editor.

You can associate domain models with bindings to restrict the values WebStudio users can enter. Add the domain models first.

Some familiarity with rules is assumed. See [Chapter 16, Rules and Rule Functions, on page 251](#) for details.

To Add a Rule Template

1. Right-click the folder where you want to store the rule template, and select **New > Other > TIBCO BusinessEvents > Rule Template**. You see the New Rule Template Wizard.
2. In the **Rule Template Name** field, type a name for the rule template. In the **Description** field, type a description.



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

3. Click **Finish**. You see the Rule Template editor.

Configuration

4. Complete the Configuration section as you would for a rule (see [Adding a Rule on page 254](#)).

Declarations and Variables

5. In the Declarations and Variables section (equivalent to the `declare` statements in the source editor), do the following:
 - a. Click **Add**. You see the Select Rule Declaration Arguments dialog box.
 - b. In the upper half of the dialog box, select the type you want to use.
 - c. If you choose an entity type, select an entity of that type, from the lower half of the dialog.
 - d. Click **OK**.
 - e. Your selection appears in the list. TIBCO BusinessEvents assigns an alias to the selected type. As desired, edit the alias.
 - f. For primitives you can also add an expression to define the initial value.

Add more types as desired.

- Bindings
6. Click the **Bindings** tab. In the Bindings section, click **Add**, select a data type, and click **OK**. A new entry appears in the Bindings list. Complete the entry as desired. You can do any of the following:
 - Replace the generated alias.
 - Define an expression to use as the initial value.
 - Specify a previously created domain model to limit what WebStudio users can enter. See [Rule Template Views on page 283](#) for an explanation.

Add more bindings as desired.

- Pre-conditions and Action Context
7. Return to the Form or Source tab and use the bindings you created as follows:
 - a. In the Pre-conditions section, write condition statements (in the TIBCO BusinessEvents rule language).

Adding pre-conditions is similar to adding conditions for a rule, except that you can add non-conditional statements in a pre-condition. See [Pre-conditions Section on page 280](#).
 - b. In the Action Context section, (in the TIBCO BusinessEvents rule language) add action statements.

Adding Action Context statements is similar to adding actions for a rule, except that only those action context statements used in business rules are actually executed.
 8. Save the resource.

Adding a Rule Template View

First add a rule template. One rule template can be associated with multiple views. See [Adding a Rule Template on page 284](#).

To Add a Rule Template View

1. Right-click the folder where you want to store the rule template view, and select **New > Other > TIBCO BusinessEvents > Rule Template View**. You see the New Rule Template View Wizard.
2. In the **Rule Template** field, click **Browse** and select the rule template whose bindings you want to use in this view.
3. In the **Rule Template View Name** field, type a name for the rule template view. In the **Description** field, type a description.



You cannot change the name in the editor. To change the name of any project element, right-click the element in BusinessEvents Studio Explorer and select **Refactor > Rename**. See [Chapter 3, Element Refactoring Operations, on page 43](#) for more details.

4. Click **Finish**. You see the Rule Template View editor.

Configuration

5. Add or edit the description.
6. As needed, browse and select a different rule template
7. As needed, click the Rule Template link to edit the selected rule template, for example to add new bindings.

Bindings defined in the rule template appear in the Bindings section of the Presentation section.

Presentation

8. In the Presentation section, use plain text or HTML to define an input form. Drag and drop the bindings into the Presentation section, or type them. The syntax for a binding is as follows (the closing tag is required):

```
<binding id="bindingName"></binding>
```

You can click the **Preview** tab to see how the view will appear in WebStudio.

9. Save the resource.

Rule Template Editor Reference



The rule template, rule, and rule function editors are similar. This section focuses on the form-based editor. You can adapt the information to apply to the different blocks of code in the source editor.

Property	Description
Configuration Section	
Same as the Rule Editor Configuration section. See Rule Editor Reference on page 258 for details.	
Declarations and Variables Section	
Type	<p>Types you add define the scope. As in the Term section of a rule, you can specify concept or event types in the project that you will use in your rule.</p> <p>In addition you can specify variables using primitive types.</p> <p>For example (in the source view):</p> <pre>Concepts.Accounts.CheckAccount checkacc; int i = 123;</pre> <p>It is not necessary to add scorecards in the declaration in order to use them in the rule or rule template.</p>
Alias	<p>A name used to refer to the scope element or variable in the Pre-conditions (when) and Action Context (actionContext) sections of the rule. Aliases must be valid identifiers See Identifier Naming Requirements on page 323.</p>
Expression	<p>If you specify a primitive you can also specify an expression. See Declarations and Variables Section on page 280 for an overview.</p>
Pre-conditions Section	
<p>Each line in the Pre-conditions area is a single expression. Expressions can be local variable declarations, method calls and so on, as well as expressions that evaluates to true or false such as are found in regular rules. Each line is joined to the others with an implicit and an operator.</p> <p>For the OR operator, use a double pipe () on the same line.</p> <p>TIBCO BusinessEvents evaluates single conditions from left to right. TIBCO BusinessEvents optimizes the evaluation of multiple conditions (see Order of Evaluation of Rule Conditions in <i>TIBCO BusinessEvents Architect's Guide</i>.)</p> <p>See Pre-conditions Section on page 280 for an overview.</p>	

Property	Description
Action Context Section	
	<p>Defines the list of all possible actions that can be executed when a rule instance based on this template is fired. Only the subset of these action context statements that are defined as commands in the business rules, however, are actually considered.</p>
	<p>When you add an action context statement, you first select one of the action types, then select an entity (create or modify action types) or rule function (call action type). See Action Context Section on page 281 for an overview.</p>

Chapter 18 **Functions**

This chapter explains how to work with functions in TIBCO BusinessEvents.

Topics

- [Overview of Catalog Functions, page 290](#)
- [Built-in Functions, page 291](#)
- [Custom Functions, page 295](#)
- [Ontology Functions, page 299](#)
- [Enabling Extended Functions, page 300](#)
- [Function Tooltips and Decorations, page 301](#)
- [Temporal Functions and Their Parameters, page 303](#)
- [Virtual Rule Functions and VRF Catalog Functions, page 305](#)
- [Cache Related Functions, page 309](#)
- [Indexing for More Efficient Cache Queries, page 311](#)
- [Oracle Coherence Cache Query Functions, page 315](#)
- [Structure of a Function Catalog, page 317](#)

Overview of Catalog Functions

The functions registry includes various catalogs of functions provided with the product, and each catalog organizes its functions into categories. You can use functions in rule conditions and actions and in rule function bodies.

Some functions work together. For example, the Standard catalog function `String.append` requires that you first use the function `String.createBuffer`, to create a buffer to which strings can be appended.

To View the Catalog functions

All catalogs appear in the Catalog Functions view. To open the view navigate to `Window > Show View > Other > TIBCO BusinessEvents` and select `Catalog Functions`. The catalog view appears on the right, by default.

To View Documentation for Functions

Documentation for functions is provided in tooltips you can access while using the Catalog Functions view. Hover the mouse over the function name to see the tooltips.

You can also access this documentation using the online functions reference, available in the HTML and Eclipse versions of the documentation, but not in PDF. To access the HTML version of documentation, open this page:

`BE_HOME/doc/standard/html/index.htm`

Expand the contents panel to `Online References > Online References`, and on the right select the *TIBCO BusinessEvents Functions Reference* link.

Built-in Functions

For all the built-in functions, this section lists the main categories in each function catalog (but not sub-categories). See [TIBCO BusinessEvents Functions Reference](#) for full details.

Standard Functions

The most generally useful catalog, the standard function catalog include the following categories:

- **Channel** functions return information about destinations, and can resume and suspend a destination.
- **Cluster** functions help with multi-engine functionality
 - **DataGrid** functions are for use with Cache object management. See [Cache Related Functions on page 309](#).
- **Collections** functions allow you to deal with various collections structures.
 - **Comparator** functions
 - **Iterator** functions
 - **List** functions
 - **Map** functions
 - **Set** functions
 - **SortedMap** functions
 - **SortedSet** functions
- **Date** functions allow you to compare two DateTime values using only the date portion of the value.
- **DateTime** functions allow you to perform these date/time related tasks and more: add units of time to a DateTime, compare, retrieve, and format dates and times.

- **Engine** functions allow you to retrieve information about the engine, for example, available memory or the number of rules fired.
 - **Locale** functions
 - **Profiler** functions
 - **Rtc** functions
 - **Variable** functions
- **Event** functions allow you to assert, create, and send simple events and perform other event-related tasks, for example, return the default destination URI of a simple event.
 - **Ext** functions
- The **Exception** function enables you to create an exception.
- **File** functions provide various useful functions used when working with files.
- **HTTP** functions are used with the HTTP channel
 - **Servlet** functions
- **Instance** functions allow you to create and delete concept instances and perform other instance-related tasks, for example, return an instance given an internal ID.
 - **PropertyArray** functions
 - **PropertyAtom** functions
 - **Reflection** functions
 - **StateMachine** functions
- **Log** functions allow you to write statements in logs.
- **Math** functions allow you to perform advanced mathematical operations.
- **Number** functions allow you to perform type conversions from and to numbers and return the maximum and minimum values for a numeric type.
- **SOAP** functions enable you to work with SOAP messages sent through an HTTP channel.
- **String** functions allow you to perform comparisons, searches, conversions, and other operations with strings.
- **System** functions allow you to send messages to a debug log, retrieve global variables, retrieve system properties, and write data to a file.
 - **ID** functions
 - **IO** functions allow the writing and closing of specific files.

- **Temporal** functions allow you to examine and perform calculations on values stored in a property's history. For information about using temporal functions, see [Temporal Functions and Their Parameters on page 303](#).
 - **Calculus** functions
 - **History** functions
 - **Numeric** functions
 - **Statistic** functions
- **Util** functions category has one sub-category for working with HashMaps.
- **VRF** functions (that is, Virtual Rule Function functions) allow you to work with decision tables. See [Virtual Rule Functions and VRF Catalog Functions on page 305](#) and see *TIBCO BusinessEvents Decision Manager User's Guide* for details.
- **XPath** functions allow you to evaluate XPath expressions.

ActiveSpaces Functions

These functions allow you to interact with ActiveSpaces metaspaces. They are used with the ActiveSpaces channels to perform operations on the ActiveSpaces metaspace and spaces connected by the channel.

Process Orchestration Functions

These functions allow you to interact with process jobs. See TIBCO BusinessEvents Process Orchestration documentation for details.

BusinessWorks Functions

BusinessWorks functions are used in ActiveMatrix BusinessWorks integration projects. See [Chapter 38, ActiveMatrix BusinessWorks Integration, on page 637](#).

CEP Load Balancer Functions

Load balancer functions are used to configure and work with the load balancer. There are two types of load balancer functions:

- **Receiver** — Allow you to create or discard a receiver, and retrieve the receiver's membership details.
- **Router** — Allow you to create or discard a router, and send an event to a remote receiver.

CEP Pattern Functions

Pattern functions are used with the pattern matcher language for identifying patterns in events. See TIBCO BusinessEvents Event Stream Processing add-on documentation for details.

CEP Query Functions

Query functions are used with the query language for querying data in the cache. See TIBCO BusinessEvents Event Stream Processing add-on documentation for details.

Communication Functions

Communication functions provide a set of catalog functions that enables TCP communication. You can create a local TCP server and a TCP client so that TIBCO BusinessEvents can communicate with data sources not otherwise available through channels, using TCP. See [Communicating with Other Sources using TCP on page 70](#) for details.

Hawk Functions

Functions to interact with Hawk micro-agents methods.

DBMS Functions

Database functions are provided for working with database concepts. See TIBCO BusinessEvents Data Modeling for more on database concepts.

Security Functions

These functions are used internally by the TIBCO BusinessEvents Decision Manager add on, for authentication.

Studio Functions

These functions allow you to use the Studio Util functions to build classes and EAR files.

Custom Functions

You can also create custom functions. Custom functions appear in the Custom Function catalog.

Adding Custom Functions

TIBCO BusinessEvents allows you to write your own custom functions in Java and add them to the function registry. When you have added the custom functions to TIBCO BusinessEvents Studio, their catalog appears automatically in the Catalog Functions view, along with the built-in function catalogs.

TIBCO BusinessEvents documentation does not contain detailed instructions for creating custom functions or programming in Java. It is assumed that you are already familiar with programming in Java, and have at a minimum already implemented a class and a static function.

See also:

- [Structure of a Function Catalog on page 317](#)
- [Example Custom Function with Tooltips on page 297](#)
- [Structure of a Function Catalog on page 317](#)

Task Summary

The steps below summarize the tasks required to integrate your own custom functions with TIBCO BusinessEvents Studio. An additional action is required to make the functions available at runtime.

1. Write a custom static function in Java and compile it. Add help content as a Java annotation, following the structure shown in [Example Custom Function with Tooltips on page 297](#).
2. Create a file called `functions.catalog`, an XML file that makes it possible to access your custom functions from TIBCO BusinessEvents Studio. You can also include information for a tool tip for each function. See [Structure of a Function Catalog on page 317](#) for details.
3. Create a `.jar` file that includes the static function class files for the implementation and the function catalog file.
4. Add the location of the JAR file or files (and any dependent JAR files) to the project properties. See [Working with External Library and Custom Function Paths on page 8](#) for details.

The functions are then available for use in any rule editor. (It is not necessary to restart TIBCO BusinessEvents Studio before you can use your functions.)



The locations of the custom function jars are stored in the `.bepproject` configuration file, which is located at the root level of the TIBCO BusinessEvents Studio project. You can check this configuration file into a version control system so that it can be shared with other project developers.

Note that the `.bepproject` file contains other information as well, such as the location of project libraries and so on.

Restrictions on Use of Custom Functions

Note the following restrictions that pertain to using custom functions:

Static and Non-Static Functions

Custom functions must be written in Java and have public static modifiers. As a workaround, encapsulate a non-static function in a static function and compile the encapsulating class to get the `.class` file.

Return Types

TIBCO BusinessEvents custom functions support the following return types:

- Java types supported are: `Object`, `String`, `Calendar` (which displays in TIBCO BusinessEvents as `DateTime`), `Integer`, `Long`, `Double`, `Boolean`, `int`, `long`, `double`, and `boolean` (but not `byte`, `short`, `float`, or `char`).
- Entity and its subtypes are supported, such as `Concept` or `Event`.
- Arrays are supported, with the exception of multidimensional arrays.

Name Overloading

The `functions.catalog` file makes functions available for use in TIBCO BusinessEvents Studio. The structure of the file requires each function within a class to have a unique name. Because of this structure, you cannot refer to an overloaded function in `functions.catalog`.

For example, the standard Java library has several `String.valueOf()` functions overloaded for each primitive type (`String.valueOf(int i)`, `String.valueOf(double d)` and so on). However, the TIBCO BusinessEvents standard function catalog has a separate function name for each data type: `valueOfBoolean()`, `valueOfDouble()`, `valueOfInt()`, and `valueOfLong()`.

See [Structure of a Function Catalog on page 317](#), for more about the `functions.catalog` file.

Editing Custom Functions

If you need to edit custom functions, you must do so in a Java editor or project, and then re-export them to a custom function JAR. Add the JAR to the TIBCO BusinessEvents project, as explained in [Structure of a Function Catalog on page 317](#).

Example Custom Function with Tooltips

The following very simple example of a custom function illustrates the required structure of a function called `add`, with the Java annotation that displays the tooltip for the function (beginning at `@FunctionJavadoc` in the example below).

```
package com.acme.functions.string;

import com.tibco.be.model.functions.FunctionJavadoc;
import com.tibco.be.model.functions.FunctionParamDescriptor;
public class StringFunctions {

    public @FunctionJavadoc (
        name="add",
        synopsis="Returns the concatenation of two strings.",
        signature="String add (String a, String b)",
        params={
            @FunctionParamDescriptor(
                name="a",
                type="String",
                desc="param 1"),
            @FunctionParamDescriptor(
                name="b",
                type="String",
                desc="param 2")},
        freturn = @FunctionParamDescriptor
            name="",
            type="String",
            desc="The concatenation of the parameter string"),

        version="5.1.0",
        see="",
        mapper=false,
        cautions="none",
        domain={"action", "condition", "query"},
        example="<br/> " +
            "String result = concatenate(\"a\", \"b\");" +
            "<br/><br/> Result is: result contains: \"ab\"."
    )
    static String add(String a, String b) {
        return a+b; // default impl.
    }
}
```

The sample above should provide enough guidance. A few additional notes follow.

synopsis Add the description of the function in the synopsis parameter.

domain Shows where the function can be used. This text is informational only. It does not appear in the tooltip. The code controlling where the function can be used (and the decorations that appear) is in the function catalog. See `<isActionOnly>`, `<isValidInQuery>`, `<isValidInBUI>` in [Table 31, Function Catalog Elements, on page 318](#).

version Specify the version of TIBCO BusinessEvents that this function supports. Informational only.

see Provide a URL for additional information.

mapper Set to true if the function uses the XSLT Mapper feature. The Mapper decoration is added if the function catalog has a `<mapper>` element. In the function catalog file is more detailed code to support use of functions in the XSLT Mapper.

Ontology Functions

Ontology functions are generated by TIBCO BusinessEvents based on the concepts, events, and rules in your project. There are three types of ontology functions:

- Constructors — Allow you to create a simple event or concept instance.
- Time events — Allow you to create and schedule a time event. See [Chapter 10, Time Events and Scheduler Functions, on page 175](#).
- Rule functions — Allow you to invoke a rule function. See [Chapter 16, Rules and Rule Functions, on page 251](#).

The Ontology Functions area uses the same folder structure as the project (or rather, a subset of that structure).

Various ontology functions appear depending on the project and the add-on products installed:

- The ontology functions catalog lists all the entity types in a project. Each type has a function for creating an instance of that type.
- Each rule function also has an ontology function enabling it to be invoked in rules.
- When TIBCO BusinessEvents Views is used additional ontology functions are available for each entity type. See TIBCO BusinessEvents Views documentatoin for details.

Enabling Extended Functions

Extended functions (sometimes called hidden functions) may be made available by TIBCO Support to address customer-specific use cases. They are also sometimes used to make legacy features available to customer who wish to continue using them. To make them visible in the Catalog Functions view, do the following:

1. Open the following file for editing:

```
BE_HOME/studio/eclipse/configuration/studio.tra
```

2. Add the appropriate property, using the specified name. For example, the following formats are typically used:

```
TIBCO.BE.function.catalog.function-catalog-name=true  
TIBCO.CEP.modules.function.catalog.function-catalog-name=true
```

3. Save the file and restart TIBCO BusinessEvents Studio.

Function Tooltips and Decorations

Tool Tips

When you float the cursor over a function in the registry, TIBCO BusinessEvents displays the description and syntax in a tool tip next to the cursor. Tooltips also form the online reference to the function catalogs. See *TIBCO BusinessEvents Functions Reference* in the HTML documentation.

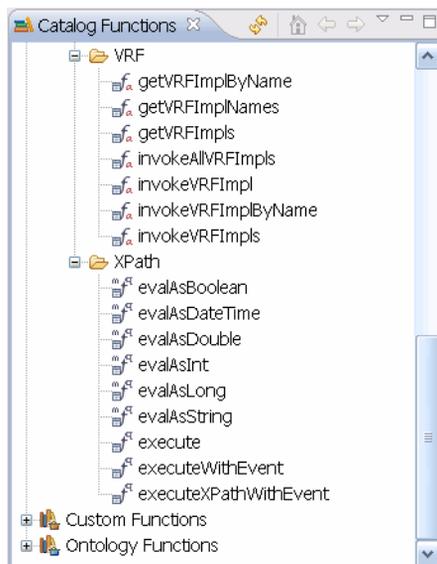


You can create your own tool tips for custom functions. See [Structure of a Function Catalog on page 317](#).

You can turn off the tool tip display too using Window >Preferences >TIBCO BusinessEvents preferences.

Decorations Indicating Where Functions can be Called

Some functions have limited application, and some can be used with various product features such as in queries (available in TIBCO BusinessEvents Event Stream Processing add-on). Such functions are decorated with small symbols that indicate useful information about use of the function. The decorations cluster around the "f" next to the function name in the function catalog. A function can have zero, one, or more decorations. The following figure shows all the available decorations. They are described in the sections after the figure:



Action-Only Functions

These functions are for use in rule actions only. Some of these functions have side effects, for example they can change values. Other functions are limited to actions for other reasons. These action-only functions are identified by a small *a* at the bottom right of the *f*. For example: *f_a setDateTime* | .

Functions that can be used both in actions and in conditions have no decoration. They are considered to have the default validity.

Mapper Functions

Functions that bring up the XSLT mapper and XPath builder are identified by a small *m* at the upper left of the *f*, for example: *^mf evalAsString* | . For more on using mapper functions, see [Using the Function Argument Mapper on page 266](#).

Functions That Can Be Used in TIBCO BusinessEvents Decision Manager

Functions that can be used in TIBCO BusinessEvents Decision Manager are marked with a small table icon, * f.getVRFImpls* | for example (which is also an action-only function).

Functions That Can Be Used in Queries or with Pattern Matcher

Functions that can be used in queries or with Pattern Matcher are marked with a blue *q* for example, * nanoTime* | . You can call such functions in a query string and in Pattern Matcher.

They are also valid in rule actions and conditions.

See documentation for TIBCO BusinessEvents Event Stream Processing add-on software for details on Pattern Matcher and query features.

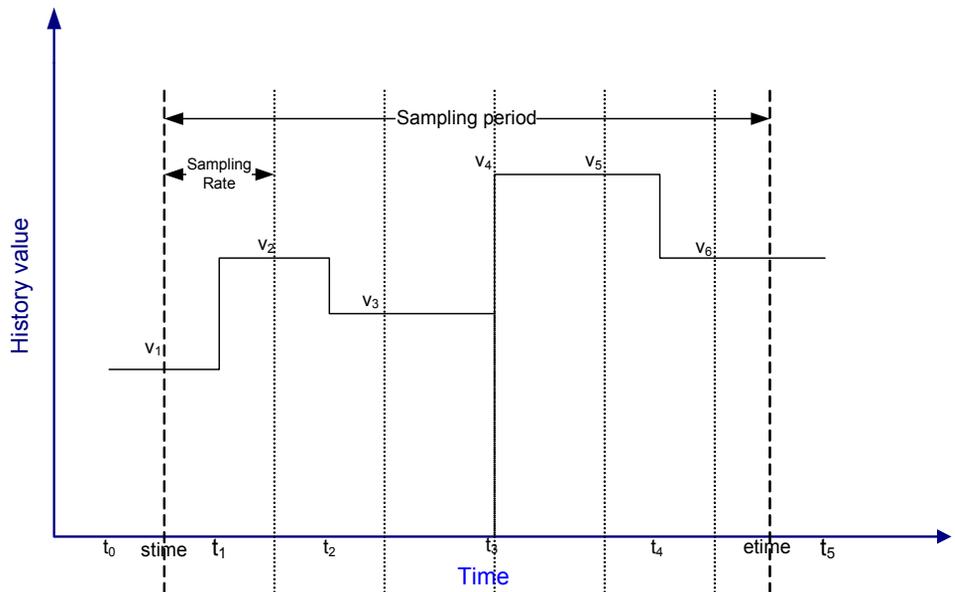
Temporal Functions and Their Parameters

The set of Built-In > Standard functions that come with TIBCO BusinessEvents includes functions that allow you to perform calculations on numeric values sampled over time. These functions are called **temporal** functions and they work exclusively with concept properties that store numeric values. Temporal functions make use of the history ring buffer to sample a property's values over time.



Use of a temporal function with a concept that has a history size of 0 may cause a runtime exception.

Figure 4 Temporal Functions Parameters



All temporal functions include these parameters, illustrated in [Figure 4](#):

- **property** — The property for which you want to sample values.
- **stime** — The time from which you want to begin sampling values (the start time) measured in milliseconds since 00:00:00 UTC on January 1, 1970.
- **etime** — The time at which you want to stop sampling values (the end time) measured in milliseconds since 00:00:00 UTC on January 1, 1970.
- **sample_rate** — The number of milliseconds between samples.

- **bound_by_stime** — A flag indicating whether the start-time is flexible:
 - **True** indicates that if the start time you provide is earlier than the timestamp for the oldest available value, you want to perform the calculation starting with the oldest available value.
 - **False** indicates that if the start time you provide is earlier than the timestamp for the oldest available value, you want to abort the calculation.

Virtual Rule Functions and VRF Catalog Functions

The VRF category of functions (within the Built-in > Standard Functions) are used only with TIBCO BusinessEvents Decision Manager.

The VRF category of functions provide flexibility when you are working with virtual rule function implementations. Virtual rule functions are implemented by decision tables.



Default Implementation When you deploy multiple implementations (tables) for one virtual rule function, but use a function that doesn't specify the implementation by name, for example if you use `Functions.MyVirtualRuleFunction()`, the default implementation is used. The default implementation is whichever was the last implementation to be deployed. However, if you use hot deployment, it may not be possible to determine which implementation was deployed last.

The VRF category of functions contains the following:

```
getVRFImplByName()
getVRFImplNames()
getVRFImpls()
invokeAllVRFImpls()
invokeVRFImpl()
invokeVRFImplByName()
invokeVRFImpls()
```

Common arguments used in the above functions are described in [Table 30, Common Arguments for VRF Functions, on page 308](#).



Defining the execution order of multiple decision tables for one VRF When a VRF has multiple implementations (decision tables), the order in which the decision tables execute can be defined using each decision table's Priority setting, which is set in the decision table Properties tab.

To Add a Virtual Rule Function

1. Right-click the folder where you want to store the virtual rule function and select **New > Rule Function**. You see the New Rule Function Wizard.
 - a. In the Rule Function Name field, type a name for the rule function.
 - b. In the Description field, type a description. (In the source editor the description appears in the * @description line of the comments at the top of the editor).
 - c. Check the **Virtual** checkbox.



You cannot change a new resource name after you click Finish. (You can change the description, however.)

2. Click **Finish**.

Click the tabs at the bottom of the editor to switch between the form-based editor and the source editor as you work, depending on your preference. These instructions use the form editor and mention the equivalent settings in the source editor.

3. In the Form editor Configuration section, add or edit a description as desired. (In the source editor the description appears in the * @description line of the comments at the top of the editor).
4. If you did not do so in the Wizard, check the Virtual check box.



In the source editor, the signature of a virtual rule function is:

```
virtual void rulefunction folder.RFName
```

Do not add code to the Body block in the source editor of a virtual rule function. If you do, you see error messages if you try to save or to switch to the form-based editor.

In virtual rule functions, the Validity field is set to Action and the Return Type is set to Void. column, select where the rule function can be used (source editor equivalents shown in parentheses):

5. In the Scope section (scope statements in the source editor), drag an ontology entity into the Scope area, OR do the following:
 - a. Click **Add** to add resources that you will be using in your rule function. You see the Select Rule Function Scope Arguments dialog.
 - b. In the upper half of the Select Rule Function Scope Arguments dialog, select the type you want to use.
 - c. If the type you select is an ontology type, in the lower half of the dialog, select a resource from the filtered ontology tree.
 - d. Click **OK**.

Your selection appears in the Declarations list. TIBCO BusinessEvents assigns an alias to it. You can edit the alias.

Add more entities as needed.

6. Save the project.

For example, here is the source view for a simple virtual rule function:

```

/**
 * @description Action to take when account is suspended
 */
virtual void rulefunction Rules.FollowUp {
    attribute {
        validity = ACTION;
    }
    scope {
        Concepts.Account account;
    }
    body {

    }
}

```

VRF Function Arguments

The VRF functions use various subsets of the following common arguments:

Table 30 Common Arguments for VRF Functions

Name	Type	Notes
vrfURI	String	The universal resource identifier (URI) for the virtual rule function. This is typically the full path to the virtual rule function within the project directory. For example, in the <code>CreditCardApplication</code> example, the virtual rule function <code>Person_VirtualRuleFunction()</code> has the following URI: <code>/Virtual_RF/Person_VirtualRuleFunction</code>
vrfImpl	Object	An object representing a virtual rule function implementation. This argument is required when invoking specific virtual rule function implementations.
implName	String	The name of a decision table (also known as a virtual rule function implementation). For example, in the <code>CreditCardApplication</code> example, the virtual rule function <code>BankUser_VirtualRuleFunction</code> has an implementation (decision table) called <code>bankUser</code> . The <code>implName</code> argument is used to retrieve a corresponding implementation object, or to execute an implementation.
args	Object array	The arguments to be passed to one or more virtual rule function implementations on invocation. These objects consist of the concepts, events, scorecards, and so on. that are needed by the implementation or implementations. For example, the <code>processApplication</code> implementation in the <code>CreditCardApplication</code> example project requires concepts of type <code>Application</code> , <code>BankUser</code> , and <code>CreditCardApplication</code> to be passed as arguments. In order to invoke the <code>processApplication</code> implementation, an instance of each concept type must be passed in the <code>args</code> array.
returnValues	Object array	<i>Not used in this release.</i> This argument is used only for the <code>invokeVRFImpls</code> function. When invoking multiple implementations, the return value of each implementation is stored in this array. The array will contain a null entry for each implementation that does not return a value.

Cache Related Functions



TIBCO BusinessEvents Express This section relates to Cache OM functionality and is not relevant if you are using TIBCO BusinessEvents Express edition.

Various standard functions in the Standard catalog `Cluster.DataGrid` category enable you to work with objects in the cache. Cache load functions load items into the Rete network so they are available.

Tool tips in the user interface (and reproduced in the *TIBCO BusinessEvents Functions Reference*) explain the details of how to use the functions. This section explains certain patterns of use.

These core functions are used with both the Oracle and TIBCO cache providers.



The previous versions of these core functions (which began with `C_`) are deprecated. Migrate projects created in earlier releases to use the current names (see [Migrating Core Coherence Functions on page 50](#)).

Query functions are not available for the TIBCO BusinessEvents DataGrid provider. Similar functionality can be achieved using the TIBCO BusinessEvents Event Stream Processing add-on.

Functions for Loading Entities to Rete from Cache and Backing Store

```
CacheLoadConceptByExtId
CacheLoadConceptsByExtId
CacheLoadConceptById
CacheLoadConceptsById
CacheLoadEventByExtId
CacheLoadEventById()
CacheLoadParent()
```

`CacheLoad*()` functions are required for working with cache only cache mode. They load entities from the cache (or backing store if not found in the cache) into the Rete network. See Loading Cache Only Objects into the Rete Network in *TIBCO BusinessEvents Architect's Guide* for details.

Use the `CacheLoad*()` functions in an event preprocessor. Only use them in rules for cases where the `ID` or `ExtId` is not known in advance (in the preprocessor).

After cache-only concepts are loaded in this way, you can then use `Instance.getByExtId()` in rules.

Never use `Instance.getByExtId()` unless you have first loaded the concept. `Instance.getByExtId()` does not assert the concept but just returns it for use in rules, for example, as read-only reference data.

Locking Functions

`Lock()`
`UnLock()`

In the event preprocessor, use the `Lock()` function to prevent other threads or engines from operating on the same entity. The lock is automatically released at the end of the RTC.

Use `UnLock()` only in a preprocessor and only to handle cases where you need to release the lock immediately instead of at the end of the RTC, for example because some information is missing that would be required to go forward.

See [Using Locks to Ensure Data Integrity Within and Across Agents in TIBCO BusinessEvents Architect's Guide](#) for details on use of locks.

Indexing Function (For Coherence Clusters Only)

`Index()` creates an index on the specified property, which is useful when you run queries. However, a simpler way is to use Present in Index checkboxes in the CDD. See [Indexing for More Efficient Cache Queries on page 311](#) for both methods.

BQL Query-Related Functions

See [To Query the Cache Using BQL Queries on page 313](#) for details

Specialty Functions Not for General Use

The following functions are used only in certain applications:

`CacheLoadConceptIndexedByExtId()`
`CacheLoadEntity()`
`CacheReevaluate()`
`EnableCacheUpdate()`

Indexing for More Efficient Cache Queries



TIBCO BusinessEvents Express This section relates to Cache OM functionality and is not relevant if you are using TIBCO BusinessEvents Express edition.

When you use Oracle Coherence as the cache provider, you can index concept and event properties to make searches faster. You can index more than one of an entity type's properties.

The query optimization you set up is used by `C_Query*()` functions, and by snapshot queries, available in TIBCO BusinessEvents Event Stream Processing add-on.

You can create the indexes in the following ways.

To Create Indexes Using a Coherence Function

This method applies to the Coherence cache provider only. It is not the preferred method. It is recommended that you use the method explained in [To Create an Index Using a Domain Object Override Setting on page 311](#).

You can create an ordered or unordered index using the following function in a startup rule function.

```
C_Index(String cacheName, Object property, boolean isOrdered)
```

where:

cacheName is a String returned by `C_CacheName()`.

property is the object returned by the appropriate `C_DatatypeAtomGetter` functions, for example, `C_StringAtomGetter()`.

isOrdered is a Boolean: set to true to order the contents of the indexed information, and set to false if you want to use an unordered index.

For example:

```
String cacheName = Coherence.C_CacheName("/Customer");
Object getter = Coherence.Extractor.C_IntAtomGetter("age");
Coherence.C_Index(cacheName, getter, true);
```

To Create an Index Using a Domain Object Override Setting

This method applies both to the TIBCO and Coherence cache providers. You can create an unordered index in the project's Cluster Deployment Descriptor (CDD) using a domain object setting.

1. Open the project CDD in TIBCO BusinessEvents Studio and go to Cluster tab > Domain Objects > Overrides.
2. Edit or create an override entry as needed for the desired entity or entities
3. In the override entry's Properties Metadata section, check the **Present in Index** checkbox for the property you want to index.
4. Save the CDD.

BQL Cache Queries



- **TIBCO BusinessEvents Express** This section relates to Cache OM functionality and is not relevant if you are using TIBCO BusinessEvents Express edition.
- This method of querying a cache requires use of the TIBCO BusinessEvents Event Stream Processing add-on product.
- The term BQL is an acronym indicating the TIBCO BusinessEvents Event Stream Processing query language.
- This is the only method that works with TIBCO BusinessEvents DataGrid.
- For the Coherence cache provider you can also use the methods explained in the section [Oracle Coherence Cache Query Functions on page 315](#).

Enabling Explicit Tuple Format for DataGrid Fields

You can set a property in the CDD so that they are instead stored as tuples. That is, TIBCO BusinessEvents attempts to create explicit space structures for each entity type.

In Shared Nothing mode, explicit tuple format is assumed internally. You should not set this property in Shared Nothing mode.

You can still use query language queries (as explained in [To Query the Cache Using BQL Queries on page 313](#)) if entities and their properties are stored as blobs. However performance is improved if they are stored as tuples.

Only certain datatypes can be stored as tuples:

- Primitive datatypes
`int, long, boolean, string, datetime, contained, reference`
 Are stored as the equivalent TIBCO BusinessEvents DataGrid types:
`integer, long, boolean, string, calendar, long, long`

- Concepts
 - Concept property arrays are stored as blob columns
- Events
 - Payload is a byte array or blob
- State Models

If an error occurs, for example due to an unsupported datatype, the item reverts to its original blob format

To Enable Explicit Tuple Format

At the cluster level, set the following field to `true`: "Store Properties As Individual Fields".

To Query the Cache Using BQL Queries

In order to execute the queries, you must first enable a dynamic query session.

1. Add a rule function in the inference agent that contains the following function: `Query.Util.startDynamicQuerySession()`.
 Configure this as a startup rule function (in the CDD file). For more on dynamic query sessions see *TIBCO BusinessEvents Event Stream Processing Query Developer's Guide*.
2. Add another rule function. Configure it as an event preprocessor (in the CDD file).
3. In the event preprocessor function, create a query string that selects entities from the cache. You can use any valid BQL query. See [Example Rule Function for a TIBCO BusinessEvents DataGrid Cache on page 314](#) for an example.
4. Use the query string in the following function:

```
Query.Util.executeInDynamicQuerySession(queryString, null, true)
```

 The function returns a list of entities from the cache.
5. Use the list returned as needed. For example, iterate over the list and load the entities, as shown in [Example Rule Function for a TIBCO BusinessEvents DataGrid Cache on page 314](#).



You can also do cache lookups using `@id` and `@extId` attributes. See [Functions for Loading Entities to Rete from Cache and Backing Store on page 309](#).

Example Rule Function for a TIBCO BusinessEvents DataGrid Cache

The following example preprocessor rule function code shows the techniques used for querying a TIBCO BusinessEvents DataGrid cache. You can then use standard functions to work with the entities returned by the query.

```

/**
 * @description
 */
void rulefunction RuleFunctions.InputEventPreprocessor {
  attribute {
    validity = ACTION;
  }
  scope {

    Events.InputEvent inputevent;
  }
  body {
    System.debugOut("Now starting InputEventPreprocessor");

    Concepts.Misc misc = Instance.getByExtId("misc");
    String[] EXTIDS_COLLECTION =
Instance.PropertyArray.toArrayString(misc.EXTID_LIST);

    String queryString =
      "select c" +
      "\n from /Concepts/TestConcept as c" +
      "\n where c@extId in"
      + "(" +
      EXTIDS_COLLECTION[0]
      + "\", \"" + EXTIDS_COLLECTION[1]
      + "\", \"" + EXTIDS_COLLECTION[2]
      + "\", \"" + EXTIDS_COLLECTION[3]
      + "\", \"" + EXTIDS_COLLECTION[4]
      + "\"";

    System.debugOut("Executing query: \n" + queryString);
    Object resultList = Query.Util.executeInDynamicQuerySession(queryString, null,
true);
    System.debugOut("Query returned: " + Query.Util.sizeOfList(resultList) + "
rows");

    while(Query.Util.sizeOfList(resultList) > 0){
      Concepts.TestConcept c = Query.Util.removeFromList(resultList, 0);
      Cluster.DataGrid.CacheLoadEntity(c);

      System.debugOut("Now loading concept " + c + " in the pre-processor");
    }
  }
}

```

Oracle Coherence Cache Query Functions



TIBCO BusinessEvents Express This section relates to Cache OM functionality and is not relevant if you are using TIBCO BusinessEvents Express edition.

These functions work only with a Coherence cluster. For information on using Oracle Coherence with TIBCO BusinessEvents, see [Chapter 24, Cache OM and Cluster Configuration, page 407](#).

You can also query a Coherence cache using the method explained in the section [To Query the Cache Using BQL Queries on page 313](#).

Also see [Indexing for More Efficient Cache Queries, page 311](#).

Constants, Extractors, and Filters Categories

The Coherence category has Constants and Extractors functions which are used in conjunction with functions in the Filters category.

Extractor functions return values for properties of different types.

Constants functions wrap constants so they can be used in filter functions. For example if a filter checks for $X = 10$, you would first wrap 10 using `C_IntConstant()`.

Functions in the Filters category enable you to use various criteria to identify a set of objects in the cache for a query. You can pass the filter to a query function.

The `C_RuleFunction()` function allows you to specify a rule function containing a custom filter condition.

Query Category



Query category functions operate only on the cache Unlike the `C_CacheLoad*()` functions, the query functions do not look in the backing store if objects are not found in the cache.

If a backing store is used Do not use query functions that delete or modify values if a backing store is used. Instead use a query to return the IDs of the entities you want to delete and use `Instance.deleteInstance()` or `Event.consumeEvent()` as needed.

```
C_CacheInvoke()
C_CacheOnlyMode_DeleteConcepts
C_CacheOnlyMode_DeleteEntities
C_CacheOnlyMode_QueryConcepts
```

```
C_CurrentContext  
C_EntryHasNext  
C_EntryIterator  
C_EntryNextValue  
C_KeyHasNext  
C_KeyIterator  
C_KeyNextValue  
C_QueryAction  
C_QueryAndLoadConcepts  
C_QueryConcepts  
C_QueryEvents  
C_QueryEvents_Order  
C_QueryIDs
```

Query functions take various actions for a specified entity or set of entities. For example, `C_CacheInvoke()` allows you to invoke a rule function for all matching entities in the cache.

For some functions you can specify the entities by passing a filter (from the Filters category).

Tool tips in the user interface (and reproduced in the *TIBCO BusinessEvents Functions Reference*) explain how to use the functions singly or in combination to achieve the desired results.

The `C_CacheOnlyMode*` functions are for use with entities that use cache-only cache mode.

Structure of a Function Catalog

A function catalog is an XML file that conforms to the schema file `function_catalog.xsd`. This allows TIBCO BusinessEvents to integrate your custom functions with the function registry in TIBCO BusinessEvents Studio. The function catalog must be in the XML format shown below and described in [Table 31](#) to map properly to the schema.



- Name the function catalog `functions.catalog`.
- Place `functions.catalog` in the root folder of the required Java archive resource (`.jar`) file.

Example Function Catalog

This example shows two functions from the standard functions catalog as an example to follow.

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<catalog name="Standard">
  <category>
    <name>System</name>
    <function>
      <name>currentTimeMillis</name>
      <class>com.tibco.be.functions.System.SystemHelper</class>
      <args></args>
      <async>>false</async>
      <reevaluate>>true</reevaluate>
      <isValidInQuery>>true</isValidInQuery>
      <isValidInBUI>>true</isValidInBUI>
      <helpUrl/>
    </function>
  </category>
  <category>
    <name>Event</name>
    <function>
      <name>createEvent</name>
      <class>com.tibco.be.functions.event.EventHelper</class>
      <args>stylesheet, entityArray</args>
      <isActionOnly>>true</isActionOnly>
      <desc>Create a event using XSLT Mapper.
        This returns a event entity</desc>
      <async>>false</async>
      <mapper>
        <enable>>true</enable>
        <type>xslt</type>
        <inputElement>
          <xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema">
            <xsd:element name="createEvent">
              <xsd:complexType>
```

```

        <xsd:sequence>
            <xsd:element name="event" type="xsd:anyType"
                minOccurs="1" maxOccurs="1"/>
        </xsd:sequence>
    </xsd:complexType>
</xsd:element>
</xsd:schema>
</inputElement>
</mapper>
<helpurl></helpurl>
<isValidInBUI>true</isValidInBUI>
</function>
</category>
</catalog>

```

Elements in the Function Catalog

Table 31 lists and describes the elements used in the function catalog.

Table 31 Function Catalog Elements

Element Name	Sub-Elements	Description
<catalog name="name ">		The root element. Attribute: name="name" where <i>name</i> is a name you provide for this functions catalog. Example: <catalog name="custom">
<category>		This is a sub-element of <catalog>. <category> is a nesting container for a set of related functions within this function catalog.
	<name>	A name you provide for this category.
<function>		A container for the information about a single function.
	<name>	The name of the function.
	<class>	The java class that implements the function.
	<desc>	Optional. A description of the function.
	<async>	Set to true if the function executes asynchronously. Set to false if the function executes synchronously.
	<helpurl>	Not used in this release.

Table 31 Function Catalog Elements (Cont'd)

Element Name	Sub-Elements	Description
	<args>	A comma separated list of descriptive names for the function's arguments. TIBCO BusinessEvents takes the argument type from the function itself.
	<isActionOnly>	If this function has side effects, for example, if it can modify values, you can only use it in action rules. Set this parameter to <code>true</code> to alert TIBCO BusinessEvents that this function has side effects and is not valid in conditions. Otherwise, set to <code>false</code> . Valid values: <code>true</code> , <code>false</code> .
	<isValidInBui>	If this function can be used in decision tables, set this element to <code>true</code> . Otherwise, set to <code>false</code> . Valid values: <code>true</code> , <code>false</code> .
	<isValidInQuery>	If this function can be used in queries, set this element to <code>true</code> . Otherwise, set to <code>false</code> . Valid values: <code>true</code> , <code>false</code> .
	<reevaluate>	Relevant only when a function is used in a condition. See Using the Reevaluate Element on page 319 for details. Valid values: <code>true</code> , <code>false</code> .

Using the Reevaluate Element

The the <reevaluate> element of a function catalog is relevant only when a function is used in a condition. Its effect is as follows.

If set to `true` then:

- TIBCO BusinessEvents does not memorize the result of the evaluation of the condition that contains this function.
- If any of the conditions is re-evaluated, then this function is also re-evaluated.

For example, <reevaluate> is set to `true` for `currentTimeMillis()`. Given this condition:

```
stock.price > 10.0;
  currentTimeMillis() - stock.time > 600000;
```

If the condition `stock.price > 10.0;` is re-evaluated, then

`currentTimeMillis()` is also re-evaluated.

If set to false then:

- TIBCO BusinessEvents calls the function during the first evaluation and stores the result is stored and used for subsequent condition evaluations.
- TIBCO BusinessEvents Studio re-evaluates the condition only if another part of the same condition changes.

In the above stock price example, if `<reevaluate>` were set to false, then the condition would be re-evaluated only if `stock.time` changes.

Chapter 19 **Rule Language Grammar**

This chapter describes the grammar for TIBCO BusinessEvents rules.

Topics

- [Rule Language Basics, page 322](#)
- [Keywords and Other Reserved Words, page 328](#)
- [Attributes, page 329](#)
- [Working with Concept and Event Properties, page 331](#)
- [Exception Handling, page 334](#)
- [Flow Control, page 337](#)

Rule Language Basics

Whitespace

Whitespace is used to separate tokens (identifiers, keywords, literals, separators, and operators) just as it is used in any written language to separate words. Whitespace is also used to format code.

These are whitespace characters, excluding line terminators:

- the ASCII SP character, also known as "space"
- the ASCII HT character, also known as "horizontal tab"
- the ASCII FF character, also known as "form feed"

Line terminators include these characters:

- the ASCII LF character, also known as "newline"
- the ASCII CR character, also known as "return"
- the ASCII CR character followed by the ASCII LF character

Comments

Comment rules as shown:

`/* text */` The text from `"/**"` to `**/"` is ignored

`// text` The text from `"/"/` to the end of the line is ignored

Separators

The following tokens are used for separators:

- ; Statement separator for conditions and actions.
- (Expression Grouping begin, or function argument list begin.
-) Expression Grouping end or function argument list end.
- , Argument list separator.

Identifier Naming Requirements

An identifier (or *name*, to use the user interface label) is an unlimited-length sequence of letters and digits, the first of which must be a letter. Letters include uppercase and lowercase ASCII Latin letters A-Z, a-z, and the underscore (_).



Do not use the dollar sign (\$).

Identifiers are case sensitive.

Identifiers cannot have spaces (except shared resource identifiers).

Identifiers may not be the same as any literal, keyword, or other reserved word. See [Keywords and Other Reserved Words on page 328](#) and [Literals on page 325](#).

Letters and digits may be drawn from the entire Unicode character set, which supports most writing scripts in use in the world today, including the large sets for Chinese, Japanese, and Korean. This allows programmers to use identifiers in their programs that are written in their native languages.

Digits include the ASCII digits 0-9.

Two identifiers are the same only if they have the same Unicode character for each letter or digit. Note that some letters look the same even though they are different Unicode characters. For example, a representation of the letter A using \u0041 is not the same as a representation of the letter A using \u0391.

Two example identifiers: `new_order E72526 creditCheck`



Here is a more succinct way for programmers to understand the requirements:

```
<Identifier> := [ <ID_START> { <ID_PART> }* ]
<ID_START> := except '$', any character for which
java.lang.Character.isJavaIdentifierStart() returns true
<ID_PART> := except '$', any character for which
java.lang.Character.isJavaIdentifierPart() returns true
```

Local Variables

You can use local variables of the following types in rule actions and rule functions:

- Primitives
- Concepts
- Events

Primitive Arrays

You can also use primitive arrays, which are fixed length. Here are examples of array declaration, initialization, and array creation expressions:

- Array declaration and initialization:

```
int i;                // int
int[] ii = {1,2,i};  // array of int
```
- Array creation with initialization expression:

```
ii = int[] {1,2,3};
```
- Array creation without initialization expression:

```
int[] arr = int[5] {};  
arr = int[5]{};
```
- Getting the length of the array:

```
int arrLength = arr@length;
```

Literals

The TIBCO BusinessEvents rule Language supports these literals:

- **int** — One or more digits without a decimal. May be positive or negative.
Examples: `4` `45` `-321` `787878`
- **long** — An integer literal suffixed with the letter L. The suffixed L can be either upper or lower case, but keep in mind that the lower case L (l) can be difficult to distinguish from the number one (1).
Examples: `01` `0777L` `0x100000000L` `2147483648L` `0xC0B0L`
- **double** — A number that can represent fractional values. D suffix is optional unless there is no decimal point or exponent.
Examples: `1D` `1e1` `2.` `.3` `0.0D` `3.14` `1e-9d` `1e137`
- **String** — Zero or more characters enclosed in double quotes (""). The string must be on one line. Use `\n` for newlines. Use the plus sign (+) to concatenate string segments.
Examples: empty string: `""` (quotes with no space). Space character: `" "` (quotes with one or more spaces). Strings with values: `"P0QSTN3"` `"The quick brown fox had quite a feast!"` Strings spanning multiple lines:
`"The quick brown fox "` +
`"had quite a feast!"`
- **boolean** — One of these two values: `true` `false`
- **Null** — This value: `null`

Escape Sequences

You can represent characters in literals using these escape sequences:

Table 32 *Escape Sequences*

Character	Escape Sequence
backspace	\b
horizontal tab	\t
linefeed	\n
form feed	\f
carriage return	\r
double quote	\"
single quote	\'
backslash	\\

Operators

The language defines the following operators. Operators in this list that are also used in the Java language work the same as the Java operators:

Table 33 *Operators in the TIBCO BusinessEvents Rule Language*

Operator	Notes
++ --	increment, decrement
+ -	unary plus, unary minus
* \ %	multiplication, division, remainder
+ -	addition, subtraction
> < >= <=	greater than, less than, greater than or equal to, less than or equal to
=	assignment
== !=	equality, inequality (Does deep string comparison, unlike Java.)

Table 33 Operators in the TIBCO BusinessEvents Rule Language (Cont'd)

Operator	Notes
&& !	boolean AND, OR, NOT
-- += *= /= %=	combined operation and assignment. As an example, <i>expr</i> += 2; is the same as <i>expr</i> = <i>expr</i> + 2; += works on strings as well as numbers. For example, you have a variable <code>String s = "abc"</code> ; and then you use <code>s+= "def"</code> ; and so the value of <i>s</i> becomes "abcdef"
instanceof	Tests whether an object is an instance of specified type. Restricted to use with concepts and events. Example: <code>boolean b = customer instanceof USCustomer;</code>
.	property access
@	attribute access

Keywords and Other Reserved Words

Do not use the words listed in this section as identifiers, resource names, or folder names. Case sensitivity depends on context (as noted in documentation). The list also includes keywords and other reserved words for all add-on products.

\$lastmod	enum	null	union
abs	Event	object	unique
abstract	except	offset	using
accept	exists	or	validity
ACTION	extends	order	virtual
AdvisoryEvent	false	order	void
alias	final	pending_count	volatile
all	finally	package	when
and	first	policy	where
as	float	priority	while
asc	for	private	
assert	forwardChain	protected	
attribute	from	public	
avg	goto	purge	
backwardChain	group	QUERY	
between	having	rank	
body	hours	requeue	
boolean	if	return	
break	implements	rule	
by	import	rulefunction	
byte	in	scope	
case	instanceof	seconds	
catch	int	select	
char	interface	short	
class	intersect	SimpleEvent	
Concept	last	sliding	
CONDITION	is_defined	static	
const	is_undefined	strictfp	
ContainedConcept	key	String	
continue	last	sum	
count	latest	super	
date	like	switch	
DateTime	limit	synchronized	
days	lock	then	
dead	long	this	
declare	maintain	throw	
default	max	throws	
delete	milliseconds	time	
desc	min	TimeEvent	
distinct	minutes	timestamp	
do	mod	transient	
double	moveto	true	
emit	native	try	
else	new	tumbling	
entity	not	undefined	

Attributes

TIBCO BusinessEvents provides attributes that you can use in rules to access information of various kinds. The attributes are listed in [Table 34](#). Use the @ operator to access attributes.

Table 34 Attributes

Entity	Attributes	Type	Returns
SimpleEvent	@id	long	The event's unique internal ID.
	@extId	string	The event's unique external ID.
	@ttl	long	The time to live of the event as specified in the configuration. (This is not the time-to-live remaining.)
	@payload	string	The payload as a string value.
Repeating TimeEvent	@id	long	The time event's unique internal ID.
	@closure	string	null.
	@interval	long	The number of units between creation of successive time events.
	@scheduledTime	dateTime	The time scheduled for asserting into the Rete network.
	@ttl	long	0.
Rule-Based TimeEvent	@id	long	The time event's unique internal ID.
	@closure	string	A string that was specified when the event was scheduled.
	@interval	long	0.
	@scheduledTime	dateTime	The time scheduled for asserting into the Rete network.
	@ttl	long	The time to live of the event as specified when scheduling the event. (This is not the time-to-live remaining.)

Table 34 Attributes (Cont'd)

Entity	Attributes	Type	Returns
Advisory Event	id	long	The advisory event's unique internal ID
	extId	String	Null
	category	String	Broad category of advisory, for example, an exception.
	type	String	Type of advisory within the category.
	message	String	Message for the user.
Concept	@id	long	The concept instance's unique internal ID.
	@extId	string	The concept instance's unique external ID.
ContainedConcept	@id	long	The contained concept instance's unique internal ID.
	@extId	string	The contained concept instance's unique external ID.
	@parent	concept	The parent concept instance. (This is treated as a concept reference in the language.)
PropertyAtom	@isSet	boolean	True if the property value has been set. Otherwise, false.
PropertyArray	@length	int	The number of PropertyAtom entries in the array.



The internal ID is automatically generated by TIBCO BusinessEvents. You cannot set it.

Working with Concept and Event Properties

This section describes how to access concept properties and event properties using the TIBCO BusinessEvents language.



Rule dependency and rule condition re-evaluation is performed only on concept properties.

This re-evaluation is not performed on an attribute such as @parent.

Accessing a Concept Property Atom

This is the syntax for accessing a concept property atom:

```
instanceName.propertyName
```

where *instanceName* is the identifier of the concept instance, and *propertyName* is the name of the concept property that you want to access.

For example to get the current value of the cost propertyAtom:

```
int x = instanceA.cost;
```

For example, to set a value with the current system timestamp:

```
instanceA.cost = value;
```



If the history size is 0, TIBCO BusinessEvents does not record a timestamp.

Get and Set PropertyAtom Value With User-Specified Time

You can get and set PropertyAtom values as follows:

- You can specify a time and get the PropertyAtom value stored in the history at that time using one of the standard functions:

```
type Instance.PropertyAtom.gettype(PropertyAtom propertyName, \
                                     long time)
```

where *type* is the type of the PropertyAtom and *propertyName* is the name of the PropertyAtom, and *time* is the time from which you want to retrieve the value.

- You can set a value in the PropertyAtom History using one of the standard functions:

```
Instance.PropertyAtom.settype(PropertyAtom propertyName, \
                                type value, long time)
```

where *type* is the type of the PropertyAtom and the type of the new value, *propertyAtomName* is the name of the PropertyAtom, *value* is the value to store in the ring buffer, and *time* is the timestamp for the new entry.

TIBCO BusinessEvents manages these requests as follows:

- **If the ring buffer has vacancies**, TIBCO BusinessEvents inserts the new entry into the correct place based on its timestamp, shifts the older values out one place, and returns True.
- **If the ring buffer is full, and the new value has a more recent timestamp than the oldest value**, TIBCO BusinessEvents inserts the new value into the correct place, shifts older values if necessary, drops the oldest value, and returns True.
- **If the ring buffer is full, and the new value has a timestamp that is older than the oldest value in the ring buffer**, TIBCO BusinessEvents does not insert the new value into the ring buffer, and it returns False.

Working with Concept Property Arrays

Accessing a Concept Property Array

This is the syntax for accessing a concept property array:

instanceName.propertyName

where *instanceName* is the identifier of the concept instance, and *propertyName* is the name of the concept property that you want to access.

Accessing a Value in a Concept Property Array

To access a value in a property array, identify the position in the array of the value as shown:

instanceName.propertyName[indexPosition]

For example:

```
String x = instanceA.lineItem[0];
```

This gets the current value of the first property atom in the array, `lineItem`, and assigns it to the local variable, `x`.



Array index difference In the TIBCO BusinessEvents language, array indexes start from zero (0). However, in XSLT and XPath languages, they start from one (1). It's important to remember this difference when using the rule language in the rule editor, and when working in the XSLT mapper and the XPath builder.

Setting the Value for an Existing Concept Property Array Position

You can set the value of an existing position in an array. For example:

```
int[] ii = {1,2,3};
ii[2] = 1;
```

Adding a Value to a Concept Property Array

You can append a value to the end of a property array. You cannot, however, add a value to any other position in an array. This is the syntax:

```
instanceName.propertyName[indexPosition] = value
```

To use the syntax shown above you must know the index position of the end of the array. You can append a value to the end of an array without knowing the index position of the end of the array using the `@length` attribute as shown:

```
instanceName.propertyName[instanceName.propertyName@length] = value
```

Deleting Values in a Property Array

This is the syntax for deleting an array property:

```
Instance.PropertyArray.delete(instanceName.propertyName, indexPosition);
```

When history is
not tracked

Array identifier numbers are positional. When history is not tracked and you delete multiple items in an array (using a for loop), delete higher position numbers before lower position numbers to ensure the correct entries are deleted. For example, if you want to delete items in positions 1, 2, and 4, delete them in this order: 4, 2, and then 1.

The reason is as follows. The array entry that held a deleted concept is removed, reducing the array size by one, and reducing by one the index of every entry in the array at a higher index than the deleted one. If you delete entries with lower position numbers first, the remaining lower numbers are then occupied by different items. For example, if you delete the item in position 1, then the item that was in position 2 is now in position 1.

Event Property

This is the syntax for accessing an event property:

```
eventName.propertyName
```

For example:

```
String x = eventA.customer;
```

where *eventName* is the identifier of the concept instance and *propertyName* is the name of the event property that you want to access.

Exception Handling

The TIBCO BusinessEvents rule Language includes try/catch/finally blocks and has an Exception type. The try/catch/finally blocks behave like their same-name Java counterparts.



Advisory Events You can also use the special AdvisoryEvent event type to be notified of exceptions that originate in user code but that are not caught with a catch block. To use the AdvisoryEvent, click the plus sign used to add a resource to the declaration. AdvisoryEvent is always available in the list of resources. See [Chapter 11, Advisory Events, on page 187](#).

This section describes the try/catch/finally commands.

Syntax

These combinations are allowed:

- try/catch
- try/finally
- try/catch/finally

```
try    try {
        try_statements
    }

catch catch (Exception identifier) {
        catch_statements
    }

finally finally {
        finally_statements
    }
```

Examples

This section provides some examples to demonstrate use of exception handling.

try/finally Example

```
String localStatus = "default status";
try {
    //readStatus might throw an exception
    localStatus = readStatus();
} finally {
    //If readStatus() throws an exception,
    //MyScorecard.status will be set to "default status"
    //but the exception won't be caught here.
    //Otherwise MyScorecard.status will be set to the
    //return value of readStatus()
    MyScorecard.status = localStatus;
}
```

try/catch/finally Example

```
String localStatus = "default status";
try {
    //readStatus might throw an exception
    localStatus = readStatus();
} catch(Exception exp) {
    System.debugOut("readStatus() threw an exception with message"
        + exp.getMessage());
} finally {
    //If readStatus throws an exception,
    //MyScorecard.status will be set to "default status"
    //Otherwise MyScorecard.status will be set to the
    //return value of readStatus()
    MyScorecard.status = localStatus;
}
```

try/catch Example

```
String localStatus = "default status";
try {
    //readStatus might throw an exception
    localStatus = readStatus();
} catch(Exception exp) {
    System.debugOut("readStatus() threw an exception with message "
        + exp.getMessage());
}
//If readStatus throws an exception,
//MyScorecard.status will be set to "default status"
//Otherwise MyScorecard.status will be set to the
//return value of readStatus()
MyScorecard.status = localStatus;
```

Flow Control

The TIBCO BusinessEvents rule Language includes commands to perform conditional branching and iteration loops. This section describes these commands.

if/else

The `if/else` command allows you to perform different tasks based on conditions.

Syntax:

```
if(condition){  
  code_block;  
}  
else{  
  code_block;  
}
```

for

The `for` command allows you to create a loop, executing a code block until the condition you specify is false.

Syntax:

```
for(initialization ; continue condition ; incrementor){  
  code_block;  
  [break;]  
  [continue;]  
}
```

`break` allows you to break out of the loop.

`continue` allows you to stop executing the code block but continue the loop.

For example:

```
for(int i=1; i<10; i=i+1){  
  System.debugOut("Hello World!");  
}
```

This example prints "Hello World!" to `debugOut` ten times.

while

The `while` command allows you to perform one or more tasks repeatedly until a given condition becomes false.

Syntax:

```
while(condition){  
    code_block;  
    [break;]  
    [continue;]  
}
```

`break` allows you to break out of the loop.

`continue` allows you to stop executing the code block but continue the loop.

Chapter 20 **Rule Language Datatypes**

This chapter provides datatype conversion tables, information about operators and types, and information about how TIBCO BusinessEvents handles inconsistency problems with datatypes.

Topics

- [Concept Properties to XML Datatype Conversions, page 340](#)
- [Compatibility of Operators with Types, page 341](#)
- [Correcting Inconsistencies of Type, page 343](#)

Concept Properties to XML Datatype Conversions

Table 35 Concept Properties to XML Datatype Conversions

Property Type	Int	Long	Float	Double	Boolean	String	DateTime	ComplexType	@ref
int	L	L	L	L		L			
long	N	L	N	N		L			
double	N	N	N	L		L			
String	L	L	L	L	L	L	L		
boolean					L	L			
Datetime						L	L		
ContainedConcept								D	
ConceptReference									ID

N - Numeric conversion, loss of information possible (see note below).

L - Shallow copy — Copies only the current value.

D - Deep copy — Copies the entire structure of the contained concept (current values of all properties only).

ID - Basically a shallow or reference-only copy. The copy refers to the same instance of the concept.



- History is never copied.
- Data loss is possible for conversions from String to a number datatype if the string represents a very large number that would have to be clipped.
- Datatype conversion tables for events are located in the TIBCO Rendezvous and TIBCO Enterprise Message Service documentation.

Compatibility of Operators with Types

Table 36, *Operator Matrix*, defines the compatibility of operators with types.

Table 36 *Operator Matrix*

		Right Side of Operator							
		str	int	lon	dou	boo	ent	obj	dat
Left Side of Operator	str	=, +, eq, cmp, inst	+	+	+	+	+	=, +, eq, cmp, inst	+
	int	+	=, math, eq, cmp	=, math, eq, cmp	=, math, eq, cmp			=, math, eq, cmp	
	lon	+	=, math, eq, cmp	=, math, eq, cmp	=, math, eq, cmp			=, math, eq, cmp	
	dou	+	=, math, eq, cmp	=, math, eq, cmp	=, math, eq, cmp			=, math, eq, cmp	
	boo	+				=, eq		=, eq	
	ent	+					=, eq, inst	=, eq, inst	
	obj	=, +, eq, cmp, inst	=, math, eq, cmp	=, math, eq, cmp	=, math, eq, cmp	=, eq	=, eq, inst	=, eq, inst	=, eq, inst
	dat	+,						=, eq, inst	=, eq, cmp, inst

Abbreviation	Meaning and Notes
boo	Boolean.
cmp	Comparison operators: <, >, <=, >=

Abbreviation	Meaning and Notes
dat	Date/Time
dou	Double
ent	Entity. Type includes Concept, Event and Scorecard. Both operands must either be of the same type or have a subtype-supertype relationship
eq	Equality operators: ==, !=
inst	instanceof
int	Integer
lon	Long
math	Numerical operators: unary +, unary -, =, -, *, /, %
obj	Object
str	String

Correcting Inconsistencies of Type

TIBCO BusinessEvents attempts to correct inconsistencies of type whenever possible by converting expressions to the appropriate type. TIBCO BusinessEvents converts expression types in the following cases:

- An expression uses the plus sign (+) with a string operand.
- An arithmetic expression includes numbers of differing types.
- The value of an expression is assigned to a variable of a different type.
- The value of an expression is passed to a function that declares a different type.

There are some inconsistencies of type that TIBCO BusinessEvents cannot correct. For example, all expressions within conditions must be of type boolean. If an expression within a condition evaluates to anything other than boolean, it would be illogical for TIBCO BusinessEvents to convert the expression to boolean. In cases like this, TIBCO BusinessEvents returns an error at compile time.

String Operands

When an expression uses the plus sign (+) with a string operand, TIBCO BusinessEvents treats the expression as a request for concatenation rather than addition. It converts the second operand to a string and concatenates the two strings.

For example:

```
"area code: " + 650 becomes
"area code: 650"
```

Arithmetic Expressions

The following information applies to these operators:

```
* / % + - < <= > = == !=
```

When an expression uses one of the above arithmetic operators with two numbers of different numeric types, TIBCO BusinessEvents promotes one of the two operands to the numeric type of the other. It makes these promotions as follows:

- If either operand is a double, TIBCO BusinessEvents promotes the other to a double.
- Otherwise, if either operand is a long, it promotes the other to a long.

Assignment Conversion

If the value of an expression is assigned to a variable, TIBCO BusinessEvents converts the expression's type to that of the variable. This might include, for example, converting a double to an int, or converting a generic model type to a more specific model type.

Function Argument Conversion

Conversions of function arguments are handled in the same way as assignment conversions.

This chapter describes how to map variables in the scope of a rule or rule function to arguments of a function used in that rule or rule function.

See also [Chapter 22, XPath Formula Builder, on page 383](#) for related information.

See [Chapter 16, Rules and Rule Functions, on page 251](#) for information about working with rules and rule functions.

Topics

- [Overview of Mapping and Transformation, page 346](#)
- [Buttons, Menus, and Icons, page 349](#)
- [Specifying Constants, page 355](#)
- [Data Validation, page 356](#)
- [Repairing Incorrect Mappings, page 357](#)
- [Shortcuts, page 358](#)
- [Examples of Mappings, page 362](#)
- [XSLT Statements, page 378](#)

Overview of Mapping and Transformation

The Function Argument Mapper allows you to supply the data that a function expects as input.

For instructions on accessing the Function Argument Mapper in the rule editor see [Using the Function Argument Mapper on page 266](#)

Function Section

The function section, in the upper part of the dialog, shows the view-only name of the function you are working with and the editable entity path to the item whose properties and attributes you want to map to the function arguments.

Input Section

Scope Variables Panel

The scope variables panel shows the list of properties and attributes available to the function, as well as global variables defined in the project.

Function Panel

The function panel uses an Extensible Stylesheet Language Transformation (XSLT) template that specifies how scope variables should be transformed to provide the expected input. Normally, you do not need detailed knowledge of XSLT to specify a function's expected output. However, if you are familiar with XSLT and you wish to see the actual code, you can right-click on any item in the Function panel and choose Copy from the popup menu. Then open a blank text document and choose Paste. The XSLT is displayed in your text document.

You can also use your own XSLT templates to perform transformations instead of using the techniques described in this chapter. You can paste XSLT into your function input fields, or into the XPath Formula field in XPath Formula Builder. (You cannot, however, paste XSLT directly into the function argument in the rule editor.)

Mapping and Transforming Data to Function Input



Assigning an empty string ("") to a field in a mapper function results in a null string.

To map data, select an item in the Scope Variables panel, then drag and drop that item into the desired schema element in the Function panel.

Simple mappings appear in the formula area to the right of the input element after you release the mouse button. For more complex mappings, click the Edit Statement () button.

Most options in the Edit Statement dialog are straightforward. However, there are some complex scenarios that require multiple steps. Many of these situations are described in the section [Examples of Mappings on page 362](#). You may also wish to refer to [XSLT Statements on page 378](#) for a reference of XSLT statements when deciding which XSLT statement can be used to achieve the result you desire.

You can specify XPath formulas to transform an element if you need to perform more complex processing. The XPath Formula Builder allows you to easily create XPath formulas. For more advanced use of XPath, see [Chapter 22, XPath Formula Builder, on page 383](#). There are also a variety of third-party books and resources about XSLT and XPath.

The datatypes of the function's arguments display as hints. Once a mapping or formula is specified, a hint becomes an XSLT statement. See [Statements, Hints, and Errors on page 347](#) for more information about hints and statements.

Statements, Hints, and Errors

When you display the Function tab, the existing statements are examined, and any input elements that do not have a statement are displayed as hints. Hints are reminders that you can specify a statement for the input element, but they are not stored as part of the XSLT template for the function's input. Hints are displayed in italics with a light blue background. Once you specify a mapping or a formula for a hint, the input element becomes a statement. You can also drag the hint to the left past the dividing line between the panels and the hint becomes a blank statement.

Once you specify a statement in the Function panel and click OK, it becomes part of the XSLT template used to create the input data. Statements are only deleted if you manually delete them using the delete button, or if you use the Mapper Check and Repair () button to automatically fix errors. Therefore, if the input schema for the function changes, your statements may no longer be valid. See [Repairing Incorrect Mappings on page 357](#) for more information about using the Mapper Check and Repair button to fix statements in the Function panel.

Any statement or hint that has an error is displayed in red. A hint is only displayed in red if it is a required input element. All required input elements must have statements specified. The Mapper Check and Repair button can help you automatically fix some errors. See [Repairing Incorrect Mappings on page 357](#) for more information about fixing errors.

Buttons, Menus, and Icons

The Input tab contains several toolbar buttons, popup menus, and icons. This section describes the various graphical elements of the Input tab.

Toolbar and Right-Click Menu on the Input Section

The Scope Variables panel and the Function panel have several buttons for performing various functions. There is also a popup menu when you right-click on elements in each panel. [Table 37](#) describes the buttons and right-click menu items available in the panels of the Input tab.

Table 37 *Input tab toolbar buttons*

Button	Right-Click Menu	Description
Scope Variables Panel		
		Coercions. Allows you to specify a type for Scope Variables elements that are not a specific datatype. For example, a choice element can be coerced into one of the possible datatypes for the element, or an element of datatype any can be coerced into a specific datatype.
		Type Documentation. Allows you to specify or view documentation for schema elements.
	Expand	This menu item has two sub-menus: Content and All. Expand > Content expands the current element so that all elements that are currently used in a mapping are visible. Expand > All expands all sub-elements of the currently selected element.
	Show Connected	Expands the elements in the Function area to display elements that are mapped to the currently selected element or its sub-elements.
	Delete	Deletes the selected element.
	Copy	Copies the selected element. The element can be later pasted.
Function Panel		
		Shows or hides the mapping formulas for the input elements.

Table 37 Input tab toolbar buttons

Button	Right-Click Menu	Description
		Move Up. Moves the selected element up in the Function tree.
		Move Down. Moves the selected element down in the Function tree.
	Move Out	Move Out. Promotes the selected element to the next highest level in the Function tree.
	Move In	Moves the currently selected element into a new statement. This displays the Move Into New Statement dialog that allows you to choose the statement you wish to move the element into. See XSLT Statements on page 378 for more information about XPath statements.
	Delete	Deletes the mapping for the selected element. If no mapping is defined, the element itself is deleted (along with all child elements). Note Elements are predefined. Do not delete elements. Deletion of an element causes mapper validation errors due to the mismatch of the right panel's content with its schema.
		Insert. Click Insert to pop-up a New XSLT Statement dialog where you can define an XSLT statement. The statement is inserted in the function input schema on the same level of the hierarchy as the currently selected element. You can add one XSLT statement at a time with this button. The right-click menu item Statement provides a shortcut for multi-line statements, such as Choice or If. See the description of the Statement menu item below for more information. Note Elements are predefined. Do not add new elements. Doing so causes a mismatch of the right panel's contents with its schema. See XSLT Statements on page 378 for more information about XSLT statements.
		Add Child. Adds a statement for a child element to the currently selected element.

Table 37 Input tab toolbar buttons

Button	Right-Click Menu	Description
		<p>Mapper Check and Repair. Verifies the XSLT template you have created in the Function panel against the expected input. A list of errors and warnings appear and you can choose which items you wish to fix. TIBCO BusinessEvents attempts to fix simple problems such as adding missing items that are expected.</p> <p>See Repairing Incorrect Mappings on page 357 for more information.</p>
		<p>Edit Statement. Allows you to modify an XSLT statement for the element.</p> <p>See XSLT Statements on page 378 for more information about XSLT statements.</p>
	Edit	<p>XPath Formula Builder. Invokes the XPath formula builder. You can use this editor to create an XPath statement for this input element. See Chapter 22, XPath Formula Builder, on page 383 for more information about XPath and the XPath formula builder.</p>
	Expand	<p>This menu item has three sub-menus: Content, Errors, and All. Expand > Content expands the current element so that all sub-elements that have a mapping or expression are visible. Expand > Errors expands the current element so that all sub-elements that have an error in their expression are visible. Expand > All expands all sub-elements of the currently selected element.</p>
	Show Connected	<p>Expands the elements in the Scope Variables panel to display elements that are mapped to the currently selected element or its sub-elements.</p>
	Statement	<p>This menu item contains shortcuts that allow you to easily add the desired XSLT statement(s) with one menu item instead of adding the statement(s) with the Insert button. See Statement Menu Options on page 358 for a description of the sub-items of this menu.</p>
	Undo operation	<p>Rolls back the last operation performed. The name of the last operation is shown.</p>
	Redo operation	<p>Performs the last operation that was undone with the Undo menu item. The name of the last operation is shown.</p>
	Cut	<p>Deletes the selected element. The element can be later pasted to a new location.</p>

Table 37 Input tab toolbar buttons

Button	Right-Click Menu	Description
	Copy	Copies the selected element.
	Paste	Pastes the last element that was copied or cut.

Icons for Schema Element Datatypes

Schema elements also have a set of associated icons to indicate their type. [Table 38](#) describes the icons used for schema items.



You can use the Type Documentation button to obtain any available documentation on any node in the Scope Variables or function input schema trees.

Table 38 Icons for schema items

Icon	Description
	Complex element that is a container for other datatypes. This is also called a <i>branch</i> in the schema tree.
	Simple string or character value.
	Simple integer value.
	Simple decimal (floating point) number.
	Simple boolean value.

Table 38 Icons for schema items

Icon	Description
	Simple Date or Time. This can be any of the following datatypes: <ul style="list-style-type: none"> • Time • Date • Date & Time • Duration • Day • Month • Year • Month & Year • Day & Month
	Simple binary (base 64) or hex binary value.
	Represents a schema item that can be any datatype. Data in this schema element can be any datatype.
	Choice. Specifies that the actual schema element can be one of a specified set of datatypes.

Qualifier Icons

Schema elements can have additional icons that specify qualifications. The qualifier icons have different meanings depending upon where they appear. For example, a question mark icon signifies an element is optional in the Scope Variables schema or in a hint in the Function panel. However, in an XSLT statement, the question mark signifies the statement is "collapsed" and an implicit "if" statement is set, but not displayed in the Function panel.

Table 39 describes the additional qualifiers that appear next to the name of schema items.

Table 39 Additional icons for hints

Qualifier	Scope Variables or Hint	Statement
	No qualifier indicates the element is required.	N/A

Table 39 Additional icons for hints

Qualifier	Scope Variables or Hint	Statement
?	A question mark indicates an optional Item.	An implicit "if" statement is set for this statement. This occurs when you map an optional element from the Scope Variables to an optional element in the function input schema or if you specify Surround element with if test on the Content tab of the Edit Statement dialog.
*	An asterisk indicates the item repeats zero or more times.	N/A
+	A plus sign indicates the item repeats one or more times.	N/A
☐	A null sign indicates the item may be set to null.	A null sign indicates the item is explicitly set to null. You can set an element explicitly to null by clicking the Edit Statement button for the element, then checking the Set Explicit Nil field on the Content tab of the Edit Statement dialog.

Specifying Constants

For each element in the Function input schema tree, you can specify a constant. Constants can be strings or numeric values. To specify a string, enclose the string in quotes. To specify a number, type the number into the schema element's mapping field. The following illustrates specifying the string "USA" for the Country item and 94304 for the PostalCode item of a function input schema.

Element	Value
Order	
OrderId	
RequiredDate	
ShipName	
CustomerName	
Street	
City	
State	
Country	"USA"
PostalCode	94304
OrderDetails	

Constants can also be used in functions and search predicates. To learn more about complex XPath expressions that use functions and search predicates, see [Chapter 22, XPath Formula Builder, on page 383](#).

Date and Datetime Strings in Constants

In constant expressions used in bindings, datetime values are read in according to the ISO 8601 standard, as described in the XML Schema specification. For example, the value:

```
"2002-02-10T14:55:31.112-08:00"
```

is 55 minutes, 31 seconds and 112 milliseconds after 2pm on February 10th, 2002 in a timezone that is 8 hours, 0 minutes behind UTC.

If no timezone field is present, the value is interpreted in the timezone of the machine that is performing the parsing. This can lead to complications if you are processing data from different timezones, so you are encouraged to always use timezones.

When TIBCO BusinessEvents generates datetime strings UTC time is always used. The output for the example above is:

```
2002-02-10T14:55:31.112Z
```

which is the equivalent time in the UTC timezone.

Data Validation

Data passed as input to a function is validated to ensure that it conforms to its specified datatype.

Table 40 describes the validation behavior. Datatype validation listed with the prefix `xsd:` is defined in the namespace <http://www.w3.org/2001/XMLSchema>. See *XML Schema Part2: Datatypes* specification at <http://www.w3.org/TR/2004/PER-xmlschema-2-20040318/> for more information on the proper representation of these datatypes. Datatype validation listed with the prefix `xdt:` is defined in the namespace <http://www.w3.org/2003/11/xpath-datatypes>. See *Xquery 1.0 and Xpath 2.0 Functions and Operators* specification at <http://www.w3.org/TR/2003/WD-xpath-functions-20031112/> for more information on the proper representation of these datatypes.

Table 40 Datatype validation

Data Type	Validation
boolean	xsd:boolean
double	xsd:double
string	xsd:string
dateTime	xsd:dateTime
long	xsd:long
int	xsd:int

Repairing Incorrect Mappings

Any incorrect statements are displayed in red in the Function panel. Errors can occur for a number of reasons. For example,

- a required element has no statement, and therefore must be specified
- the function's input schema has changed and existing statements may no longer be valid
- the XPath formula for an element may contain an error

You should correct any errors before attempting to test or deploy your process definition.

To help find potential problems in your mappings, click the Mapper Check and Repair button. This button displays a dialog with all potential problems in the specified mappings. You can select the Fix checkbox for potential errors, and TIBCO BusinessEvents will attempt to fix the problem.

Some potential problems in the Mapper Check or Repair dialog cannot be fixed easily, and therefore there is no checkbox in the Fix column for these items. For example, if an element expects a string and you supply a complex type, the corrective action to fix the problem is not clear, and therefore TIBCO BusinessEvents cannot automatically fix the problem. You must repair these items manually.

If you want to return to the original expected Function input and remove all of the currently specified mappings, perform the following:

1. Delete the root element of the function's input by selecting it and clicking the Delete button.
2. Click the Mapper Check and Repair  button.
3. In the Mapper Check and Repair dialog Fix column, select the checkboxes for all items.
4. Click OK.

Alternatively, you can simply select the root input element and press the delete key on your keyboard as a shortcut for the procedure above.

After deleting all mappings and schema items and then repairing the input schema, the function's input reverts to the state before you open the Argument Mapping Wizard for the first time for this function.

Shortcuts

The Move In, Insert, Add Child, and Edit Statement buttons on the toolbar are ways to manually manipulate XSLT statements in the Function panel. These buttons, however, only add or modify one statement at a time. Also, there are some situations where you wish to convert a hint into a statement without performing any mapping. This section describes shortcuts for manipulating XSLT statements.

Statement Menu Options

When you select an element in the function input schema and right-click, a menu appears. The Statement menu item contains several sub-items that are useful shortcuts for creating XSLT statements.

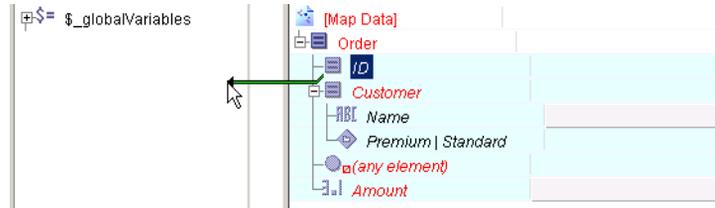
- Surround with Choice — a shortcut for adding a choice statement and its associated conditions or otherwise statements around the currently selected element.
- Surround with If — a shortcut for adding an if statement and placing the currently selected element as the sub-element of the if.
- Surround with For-Each — a shortcut for moving the current element into a For-Each statement.
- Surround with For-Each-Group — a shortcut for moving the current element into a For-Each-Group statement and adding a Group-By grouping statement.
- Duplicate — a shortcut for creating a duplicate of the currently selected element (including any mappings or XPath formulas for the element). The duplicate is added below the currently selected element.
- Insert Model Group Content — a shortcut for inserting the contents of a selected model group into the mapper tree. The selected element in the function input schema is replaced by the contents of the model group you select.

Dragging to the Left

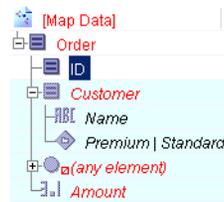
Dragging an element in the function input schema to the left past the divider between the two areas of the Input tab changes a hint into an XSLT statement. [Figure 5](#) illustrates dragging an element to the left.

Figure 5 Dragging to the left to change a hint to a statement

Before dragging, the ID element is a hint.



Dragging the ID element past the divider turns the hint into a statement.



This shortcut is useful in the following situations:

- When you have a complex element with no sub elements and no content.
- When you have a choice element, dragging to the left brings up the Mapping Wizard and allows you to choose a type for the element.
- When you have an element of type Any, dragging to the left brings up a dialog that allows you to specify the type for the element.

Cutting and Pasting

The Function panel is an XSLT template for specifying the function's input schema. You can choose any element in the Function panel and select Copy from the right-click menu or press the Control-C keys to copy the XSLT statement for the element. Once the XSLT is copied, you can paste it into a text editing tool to view or modify the code.

You can also paste arbitrary XSLT code into the Function panel using the right-click menu or the Control-V keys. Pasting XSLT code from the copy buffer places the code above the currently selected element in the Function panel.

Automatic Testing (at Runtime)

When you map Scope Variables elements to Function elements, the behavior of the mapping depends upon the types of elements you are mapping. In the simplest case of mapping a required element in the Scope Variables schema to a required Function element, the value of the Scope Variables element is assigned to the required Function element.

However, when elements are optional or nillable, more complex tests are necessary. When you drag the Scope Variables element to the Function element, the necessary tests are automatically placed into the Function XSLT template.

This section describes the result of mapping different types of elements. The types of mappings are described, then an example is given that illustrates these mappings and shows the XSLT code that is generated automatically when these mappings are performed at runtime.

Required to Required

Specifies that the statement should always include the required Function element and its value should be obtained from the required Scope Variables element that the element is mapped to.

Optional to Optional

Specifies that the statement should test if the Scope Variables element is present, and if so, include the optional element in the function's input. If the Scope Variables element is not present, the optional element is omitted from the function's input.

Nillable to Nillable

Specifies that both the Scope Variables and Function elements can be nil. Therefore, the value of the Function element is set to the value of the Scope Variables element. The value of the Function element is set explicitly to nil if that is the value of the Scope Variables element.

Optional to Nillable

Specifies that the statement should test if the optional Scope Variables element exists. If the element exists, the Function element should be created and set to the value of the Scope Variables element. If the Scope Variables element does not exist, the element is omitted from the function input schema.

Nillable to Optional

Specifies that the statement should test if the Scope Variables element has a value specified, and if so, the optional element in the Function panel should be set to the value of the Scope Variables element. Otherwise, if the Scope Variables element is nil, the optional element is omitted from the Function panel.

Optional & Nillable to Optional & Nillable

Specifies that if the optional Scope Variables element exists, then include the optional Function element in the input schema. If the Scope Variables element is nil, set the value of the Function element explicitly to nil. If the Scope Variables element is not nil, set the value of the Function element to the value of the Scope Variables element. If the Scope Variables element is not present, then omit the optional element from the function input schema.

Examples of Mappings

Some mappings require several steps to achieve the desired results. This section describes some complicated mapping scenarios and how to achieve the desired mappings using the tools available.

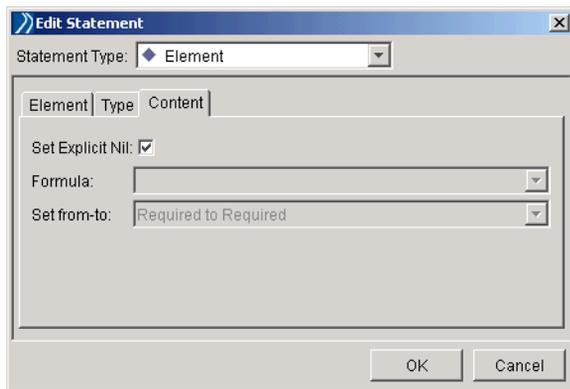


There are many methods to insert or modify XSLT statements in the function input schema. The examples in this section illustrate the simplest procedures to obtain the desired results. However, you do not have to follow the same procedures outlined in this section to achieve the correct set of statements.

Setting an Element Explicitly to Nil

In some situations, you may wish to set an element explicitly to nil. One situation is when you wish to insert a row into a database table and you wish to supply a NULL for one of the columns. To set an input element explicitly to nil, perform the following:

1. Select the input element you wish to set to nil.
2. Click the Edit Statement () button on the Input tab toolbar.
3. Click the Content tab of the Edit Statement dialog.
4. Check the checkbox in the Set Explicit Nil field.

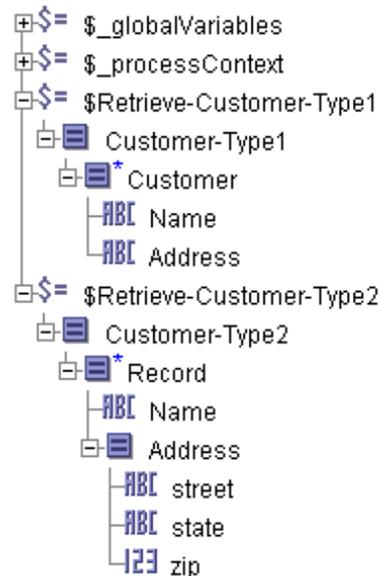


The element's formula becomes blank and is not editable (because nil is the value of the element) and the explicit nil qualifier icon appears next to the statement



Merging Input from Multiple Sources

You may have multiple items in the Scope Variables that you wish to map to one repeating element in the Function panel. For example, you may have multiple formats for customer records and you wish to create a single, merged mailing list containing all customers in one format. In this example, the schemas are the following:



You wish to combine customers into a single repeating structure to create a mailing list.

Multiple types of customers are retrieved, each having a different format for the address.

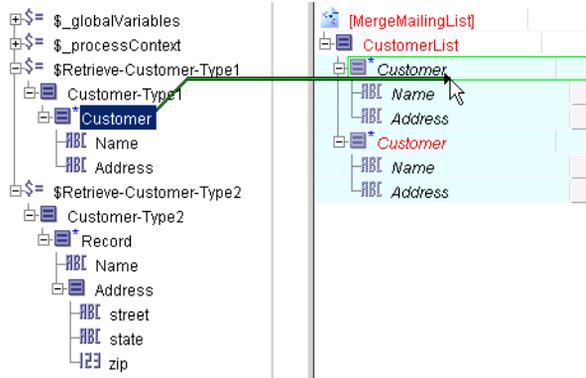
The following procedure describes how to map multiple elements into a single repeating element.

1. Select the repeating element in the Function panel, right-click, and select **Statement > Duplicate** from the menu.

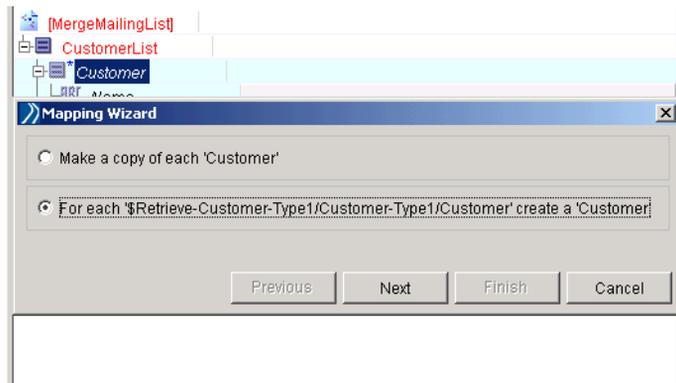
Because you are creating two different formulas for mapping, you need two copies of the repeating element, one for each format. The resulting output contains only one repeating customer element, but the two copies in the Function panel make it simpler to perform two different mappings.

2. Map one of the elements from the Scope Variables to the first copy of the repeating element in the Function. For example, map

\$Retrieve-Customer-Type1/Customer to MergeMailingList/CustomerList/Customer.



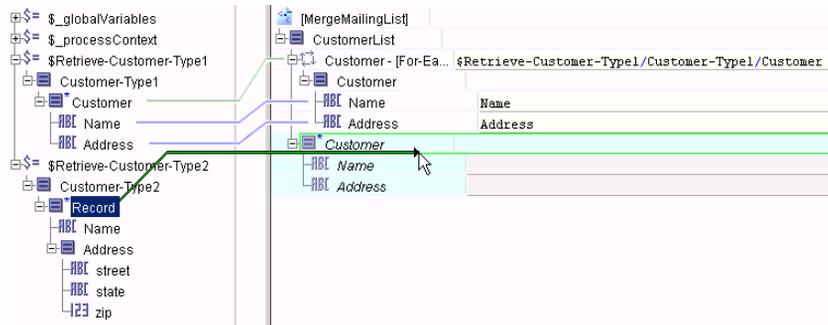
The Mapping Wizard dialog appears and presents choices for what you would like to accomplish. Choose the For Each option and click **Next**.



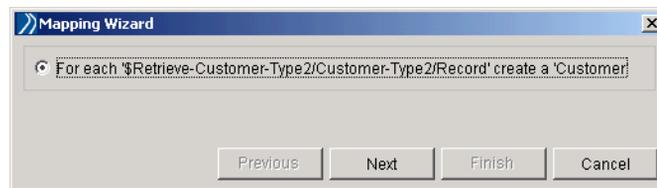
The mapping wizard asks if you wish to automatically map items with the same names. Click **Finish** to accept the default mappings.

3. Map the other element from the Scope Variables to the second copy of the repeating element in the Function. For example, map

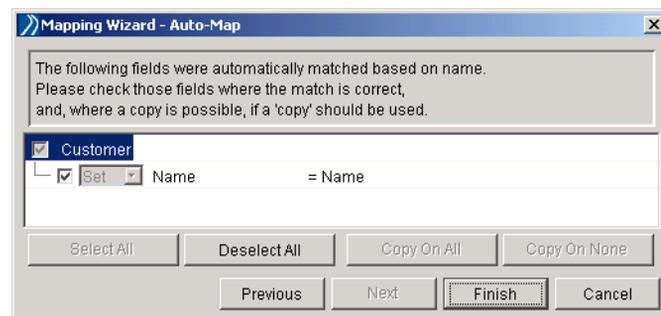
\$Retrieve-Customer-Type2/Record to MergeMailingList/CustomerList/Customer.



In the Mapping Wizard dialog, choose the For Each option and click Next.

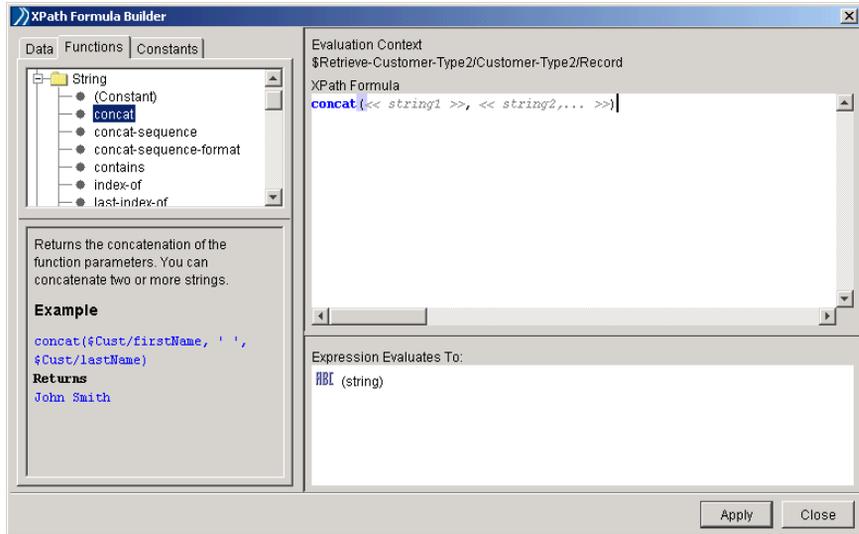


The mapping wizard presents you with an option to automatically map elements with the same name. Click Finish to accept the default mappings.

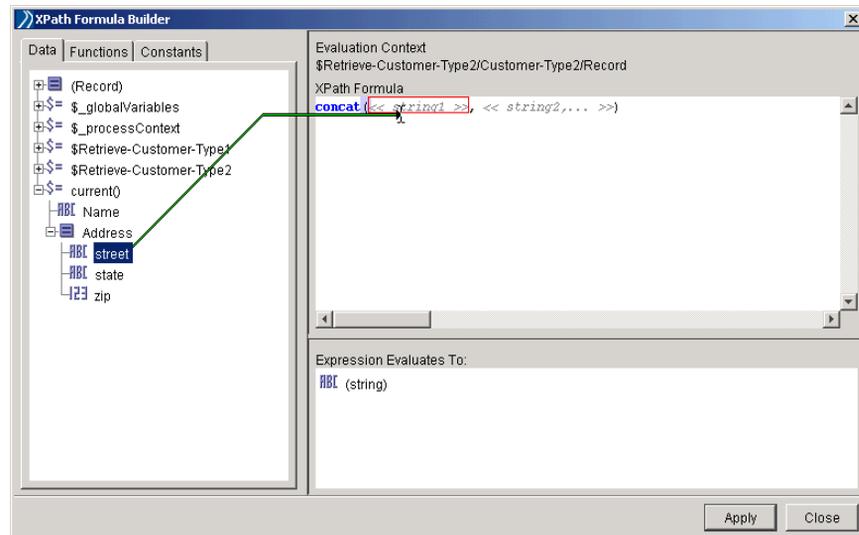


4. Select the Address element and click the XPath Formula Builder icon in the Input tab toolbar. In the XPath Formula Builder, drag a concat() function into the XPath Formula field. This function is used to concatenate the three

elements in the Record element in the Scope Variables area to one Address element in the function's input.

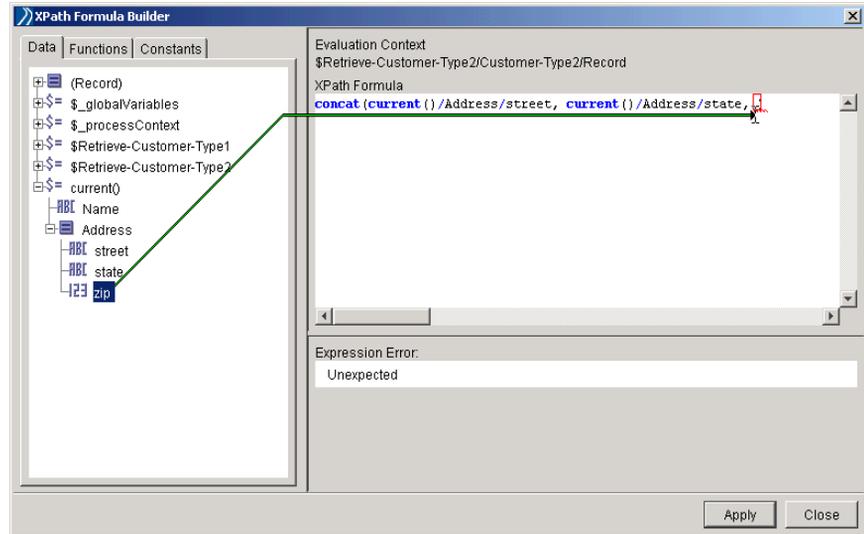


Click the Data tab, then drag the `$current()/Address/street` element into the `<< string1 >>` placeholder in the `concat()` function.

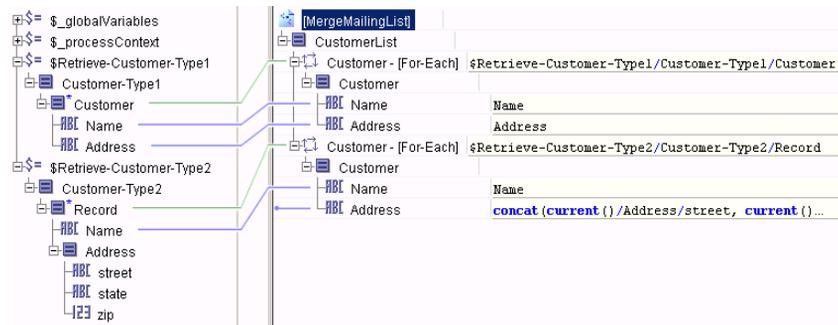


Drag the `$current()/Address/state` element into the `<< string2 >>` placeholder in the `concat()` function. Then, add a comma to the end of the function to include a third string to concatenate. Drag the

`$current()/Address/zip` element into the position of the third string in the `concat()` function.



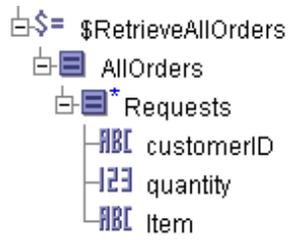
5. Click **Apply**, then click **Close** to dismiss the XPath Formula Builder dialog.
6. This results in the following mapping:



Converting a List Into a Grouped List

You may need to convert a flat list of items into a more structured list. For example, you may have a list of all orders that have been completed. You may want to organize that list so that you can group the orders placed by each customer.

In this example, the schemas are the following:



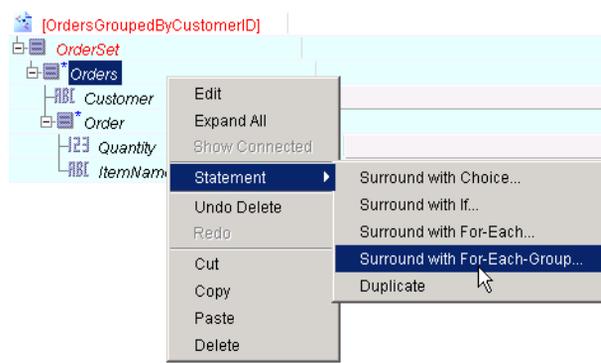
Flat list of orders placed by all customers. Each item in the repeating element Requests lists an order. The same customer can have multiple orders in this list.



The resulting schema is a repeating list of Orders, each item in the list is a customer ID. Each customer ID will, in turn, contain a list of orders.

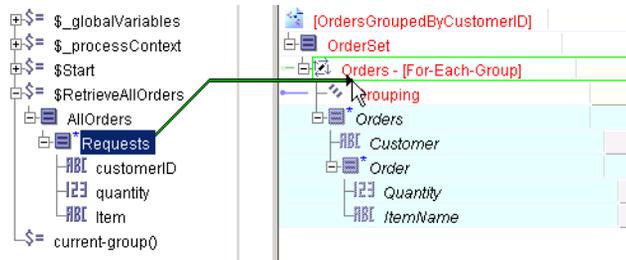
The following procedure describes how to map the flat list of orders into a list grouped by customer ID.

1. Choose the repeating element in the function input schema that holds the grouped data. In this example, that element is `Orders`. Right-click on this element and choose **Statement > Surround with For-Each-Group...** from the pop-up menu. This is a shortcut to create a For-Each-Group statement with the `Orders` element as a child element and a Grouping statement to contain the element you wish to group-by.



Adding the Grouping statement creates the `=$current-group()` element in the Scope Variables area. The Grouping statement creates the list grouped by the desired element, and the `current-group()` function allows you to access the items in the `Requests` repeating element that correspond to the group that is currently being processed.

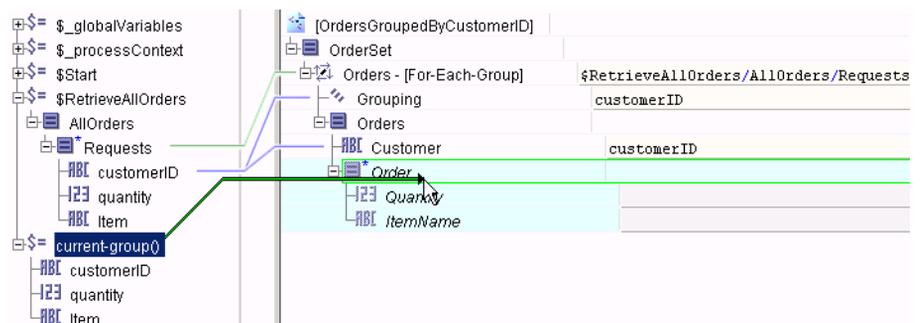
2. Drag the repeating element from the Scope Variables area to the For-Each-Group statement.



3. Drag the element you wish to group by from the Scope Variables area to the Grouping statement in the Function panel. In this example, customerID is the grouping element.



4. Map the current-group() element in the Scope Variables area to the repeating element Order under the Customer element in the Function panel.

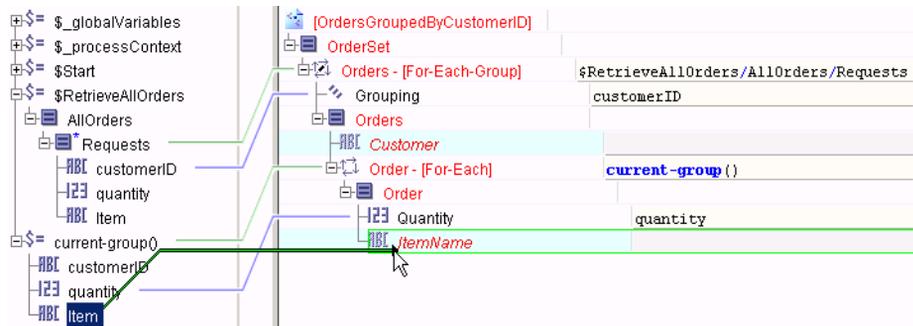


The default choice in the mapping wizard for this mapping is to create a For-Each. Choose this option in the mapping wizard.

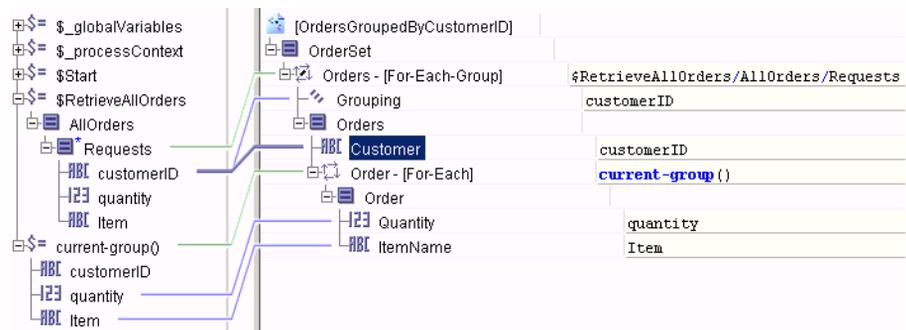


This creates an item in the Order list for each item in the current customer ID group that is being processed. The mapping wizard asks if you wish to map items with the same name in the current group and the orders group.

- Map the remaining element from the `current-group()` element into the desired element in the For-Each group. In this case, `quantity` would map to `Quantity` automatically, and `Item` must be mapped to `ItemName`.

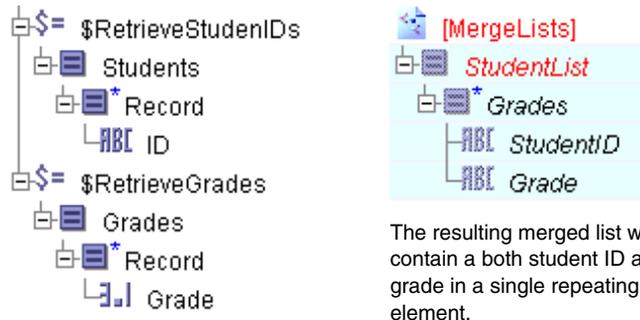


- Map the `customerID` element in the `Requests` element into the `Customer` element in the Function panel.



Merging Two Corresponding Lists

You may need to merge two lists that have corresponding items. For example, you may have a list of student IDs and a list of grades, each grade corresponds to the student ID in the same position in the student ID list. In this example, the schemas are the following:

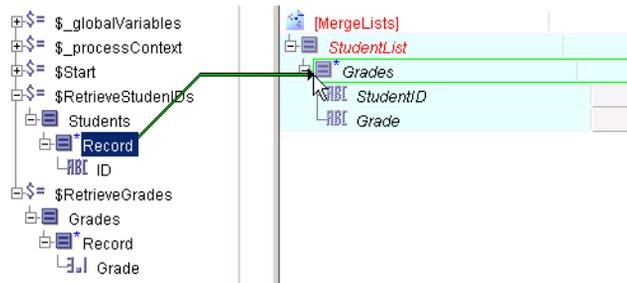


Two repeating lists contain corresponding student IDs and grades.

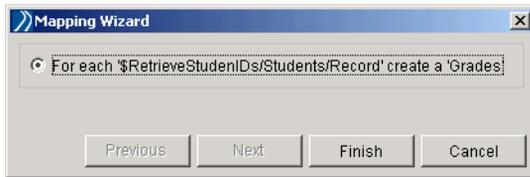
The resulting merged list will contain both student ID and grade in a single repeating element.

The following procedure describes how to merge the two repeating elements containing corresponding data into one repeating element.

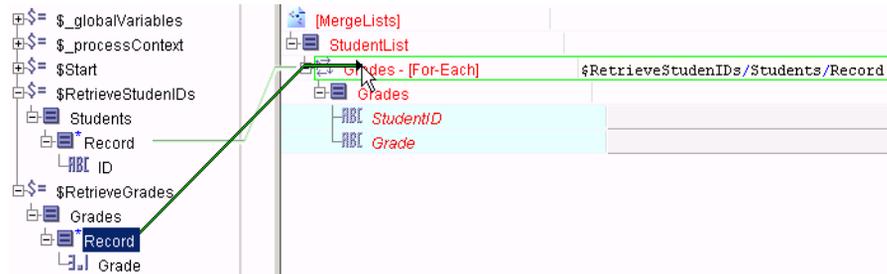
1. Map the first repeating element from the Scope Variables area into the Grades repeating element in the Function panel. In this example, the \$RetrieveStudentIDs/Students/Record is the first repeating element.



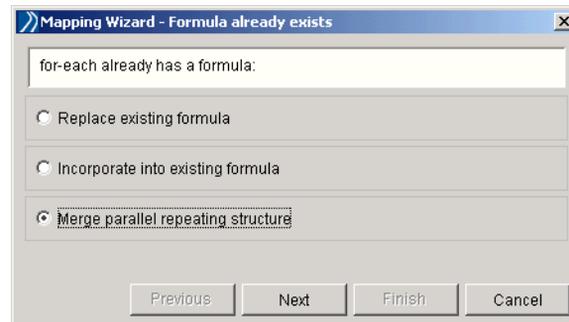
This brings up the mapping wizard with the default choice of creating a For-Each statement. Click **Finish** in the Mapping Wizard dialog to create the For-Each statement.



2. Drag the second repeating element into the For-Each statement.

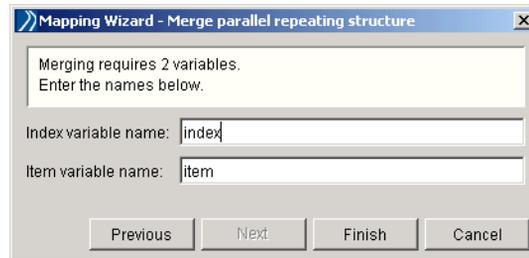


3. The Mapping Wizard dialog appears asking you to chose an option. Choose the Merge parallel repeating structure option and click **Next**.



4. Merging two parallel repeating structures requires two variables. The mapping wizard prompts you to name these two variables. One variable is to hold the position number of the current item being processed, and the other variable is to hold the item in the second list that corresponds to the position of the item in the first list. Create the variables with the default names

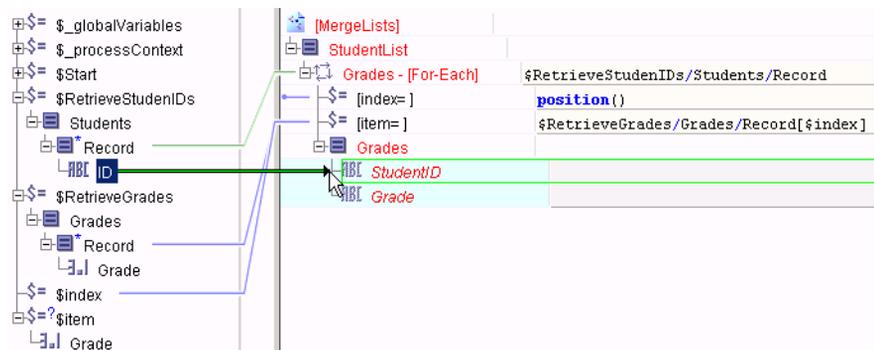
supplied by the mapping wizard, or choose your own names for these variables. Click **Finish** to proceed.



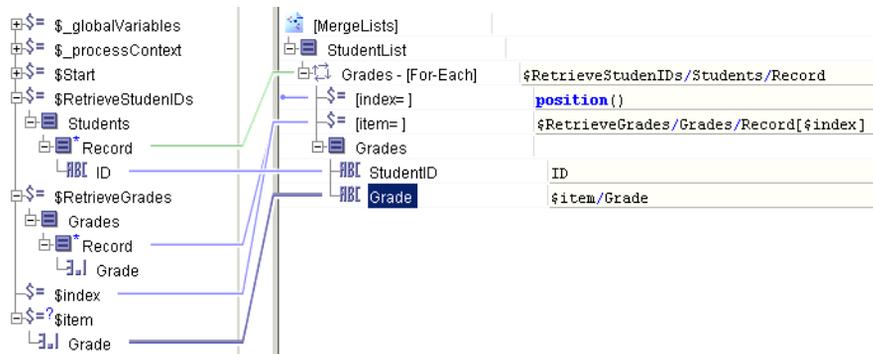
The two variables appear in the Scope Variables area once you have completed this step. The two variables also appear in the Function panel with the correct XPath statement to produce the desired result.

The `#[index=]` element contains the XPath formula `position()` to set the element with the current position number of the list item being processed. The `#[item=]` element contains a statement to retrieve the item in the second repeating element that corresponds to the position of the item in the first list that is currently being processed.

5. Map the ID element to the StudentID element in the function arguments.



6. Map the `=$item/Grade` element to the Grade element in the Function panel.

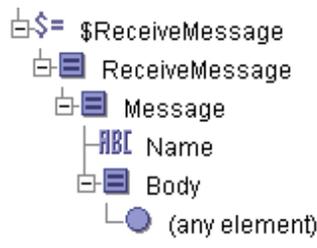


Coercions

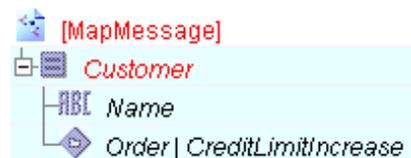
In some situations, the datatype of a Scope Variables element may be undefined. In these situations, you may know the datatype of the element, and you can coerce the element into a specific type. The Coercions button in the Input tab toolbar allows you to create and manage your coercions.

The following example illustrates a schema with an element defined as the "any element" datatype. The schema is for a generic incoming message that can have any type of body. In the example, however, the any element is coerced into an Order type so that it can be mapped to a choice element.

In this example, the schemas are the following:



The incoming message can have a body of type any element.



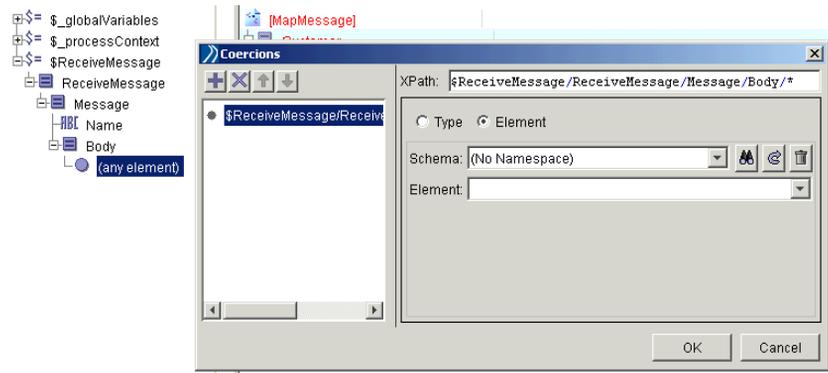
The function is expecting either an Order or a CreditLimitIncrease.

The following procedure describes how to coerce the Body element of the incoming message into a specific datatype and map it to a choice element.



There are many ways of accomplishing the same result as this example. This example attempts to illustrate the simplest method to achieve the desired result.

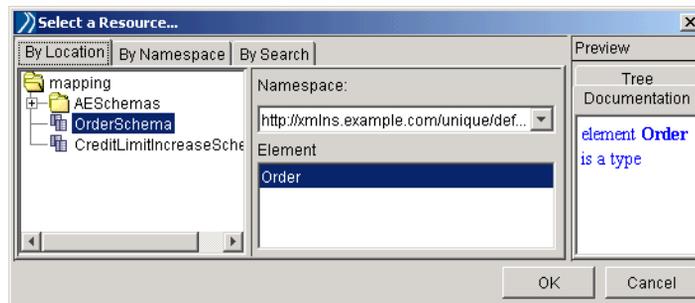
1. Select the element of type any element in the Scope Variables schema. Click the Coercions button in the Input tab toolbar. In the Coercions dialog, click the Insert button (+) to add a coercion for the currently selected element.



The Coercions dialog allows you to manage all of your coercions for a function in one dialog. You can create, modify, or delete coercions for any element in the Scope Variables schema using this dialog, not just the currently selected element. If you are creating a coercion for an element that is not currently selected, use the XPath field to specify the location of the element.

Click the Element radio button to specify that you are specifying a schema element.

2. Click the Browse Resources button next to the Schema field to browse a list of schemas that can be used. In the Select a Resource... dialog, select the schema that you would like to specify



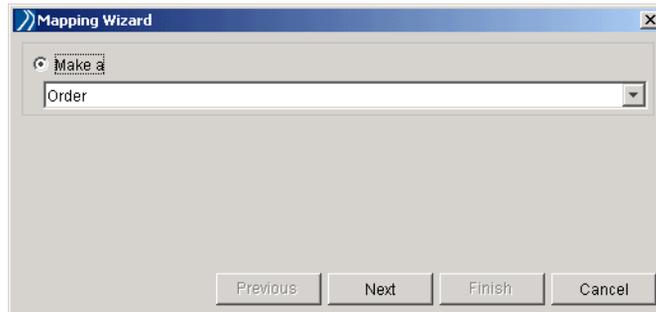
Click OK to coerce the element into the datatype of the selected schema element. The following would be the resulting schema where the element of the datatype any element has been replaced with the `Order` schema.



3. Map the Name element to the Name element in the Function panel. Then, map the coerced Order element to the choice element in the Function panel.



The Mapping Wizard dialog appears and asks if you wish to create an `Order` or a `CreditLimitIncrease` element. Select `Order` and click **Next**.

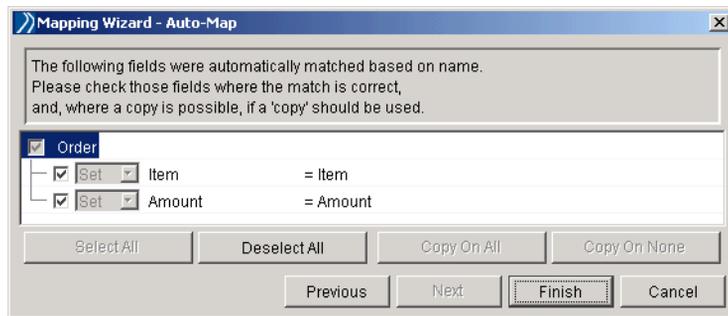


The Mapping Wizard then asks you to create a For Each. Even though there is only one element in the Scope Variables schema (the `Message` element is not

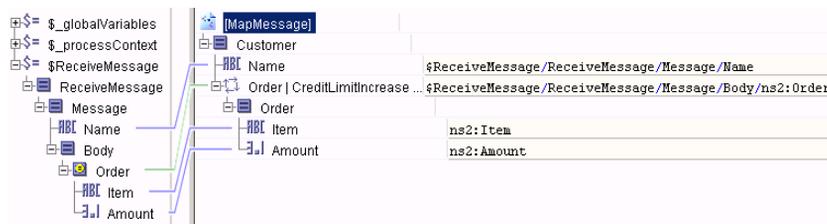
repeating), a For Each is used because this construct allows you to map the individual items of the Order element. Click **Next** to continue.



The Mapping Wizard then asks if you wish to automatically map elements with the same name. Click **Finish** to accept the default mappings.



4. The following is the completed mapping.



XSLT Statements

The following sections describe the XSLT statements you can add to your mapping. You can add or edit these statements by clicking the Edit Statement (ⓘ) button or these statements can be added automatically by selecting them from the dialogs that appear when you drag and drop elements from the Scope Variables tree to the function argument tree.

The following sections discuss statement types (available in the Statement Type drop-down list in the Edit Statement dialog).

Attribute

Allows you to specify an attribute, and optionally the namespace for the attribute. You can also specify the type of value for the attribute.

XSLT Equivalent

The following is an attribute named "lastname".

```
<ns:attribute namespace="mns" name="lastName"/>
```

When attributes are created, you can optionally specify the kind of value the attribute will have and whether the attribute should be surrounded by an if statement. For example, you can specify the value of the last name attribute to be a constant, like so:

```
<ns:attribute namespace="mns" name="lastName"/>
  "Smith"
</ns:attribute>
```

Choose

Provides a way to select transformation to perform based on an expression. Specify the condition in the when element as an XPath expression. You can optionally specify an otherwise condition for processing all elements that do not meet any of the specified when conditions.

XSLT Equivalent

The following determines if the node set for FilesTransferred contains any files, and if so, performs an action. If the node set is empty (no files were transferred), a different action is performed.

```
<ns0:choose xmlns:ns0="http://www.w3.org/1999/XSL/Transform">
  <ns0:when test="$FTP-Put/FTPputOutputFile/FileTransferred" >
```

```

    < something here ... >
  </ns:0when>
  <ns0:otherwise>
    < something here ...>
  </ns0:otherwise>
</ns0:choose>

```

Comment

Places a comment in the XSLT template. Comments are delimited by `<!--` and `-->`.

XSLT Equivalent

```
<!-- comment here -->
```

Copy

Copies the selected node to the current node in the input tree. Only the node is copied, no children of the node are copied.

XSLT Equivalent

```
<ns0:copy xmlns:ns0="http://www.w3.org/1999/XSL/Transform"
select="$Query/resultSet"/>
```

Copy-Contents-Of

Copies the selected node's contents. This is useful if you wish to copy an element to a new element with a different name.

XSLT Equivalent

```
<ns:element namespace="foo" name="bar">
  <ns:copy-of select="null/@*" />
  <ns:copy-of select="null/node()" />
</ns:element>
```

Copy-Of

Creates a copy of the selected node, including the node's children. Both the copied node and the destination node must have the same name and structure.

XSLT Equivalent

```
<ns0:copy-of xmlns:ns0="http://www.w3.org/1999/XSL/Transform"
select="" />
```

Element

Creates an element with the specified name.

XSLT Equivalent

```
<elementName>value</elementName>
```

For-Each

Performs the specified statements once for each item in the selected node. This is useful if you wish to process each item of a repeating element once.

XSLT Equivalent

The following iterates over the list of files transferred from a ActiveMatrix BusinessWorks FTP Put activity and outputs an element with the name of each file for each file transferred.

```
<ns:for-each select="$FTP-Put/FTPPutOutputFile/FileTransferred">
  <fileName>
    <ns:value-of
select="$FTP-Put/FTPPutOutputFile/FileTransferred/Name" />
  </fileName>
</ns:for-each>
```

For-Each-Group

Groups the items in a list by a specified element. This statement requires a Grouping statement to specify which element to group-by. See [Converting a List Into a Grouped List on page 367](#) for an example of using the For-Each-Group statement.

XSLT Equivalent

```
<ns0:for-each-group
xmlns:ns0="http://www.w3.org/1999/XSL/Transform" select="" />
```

Generate Comment

Places a comment element into the XSLT template. This comment will be generated into the function's output.

Comment elements have the following syntax:

```
<ns0:comment xmlns:ns0="http://www.w3.org/1999/XSL/Transform"/>
```

Generate PI

Places a processing instruction into the XSLT template.

XSLT Equivalent

```
<ns0:processing-instruction
xmlns:ns0="http://www.w3.org/1999/XSL/Transform" name="" />
```

If

An if statement is used to surround other statements in an XSLT template to perform conditional processing. If the test attribute evaluates to true, the statements in the if are output, otherwise they are not output.

XSLT Equivalent

The following if statement surrounds an attribute for processing order items.

```
<ns:if xmlns:ns="http://www.w3.org/1999/XSL/Transform"
test="not(position()=last())">
  <ns:attribute name="OrderItem">
    <ns:value-of select=
      "$GetOrderInformation/OrderInformation/OrderDetails/OrderItem"/>
    </ns:attribute>
  </ns:if>
```

Value-Of

Specifies a value-of statement. This is normally done implicitly by specifying the formula for an element (field) in the mapping, but you may insert this statement explicitly.

XSLT Equivalent

```
<ns:value-of xmlns:ns="http://www.w3.org/1999/XSL/Transform"
select="" />
```

Variable

Adds a local variable for use in the current mapping. You can specify the name of the variable and whether you wish the variable to have a select attribute.

When you add a local variable, it appears in the Function and Scope Variables panels. You can supply any XPath expression to the new variable in the Function panel (either through mapping or through the XPath Formula Builder).

Once the variable's contents have been supplied, the variable (in the Scope Variables area) can be mapped to any item.

Adding a variable is useful when you wish to join two repeating elements into a single list, then map the combined list to an item. Adding a variable is also useful if you perform the same computation repeatedly. You can map the results of the computation to several items instead of recreating the computation for each item.

Variables can also improve performance of mappings for large data structures. For example, if you have a process variable with 40 sub-elements, and you map each of the sub-elements to a corresponding input item, TIBCO BusinessEvents must retrieve the current process variable for each XPath expression, in this case 40 times. If this mapping appears in a loop, the retrieval of the current process variable occurs 40 times per iteration of the loop. With a variable, the data is retrieved only once and used for all mappings containing the variable. Therefore, to improve performance, create a local variable to hold process variables with a large number of elements and use the local variable in XPath expressions instead of the process variable.

XSLT Equivalent

```
<ns0:variable xmlns:ns0="http://www.w3.org/1999/XSL/Transform"
name="var" select="$RetrieveResults/resultSet"/>
```

Chapter 22 XPath Formula Builder

TIBCO BusinessEvents uses XPath in the XPath Formula Builder, available in the Function Argument Mapper tool (see [Chapter 21, Mapping and Transforming Data, on page 345](#)). You can use XPath, for example, when defining payloads for events TIBCO BusinessEvents also uses XPath as the language for defining conditions and transformations. This section covers the basics of XPath and its use in TIBCO BusinessEvents.

Topics

- [XPath Basics, page 384](#)
- [The XPath Formula Builder, page 387](#)
- [String Representations of Datatypes, page 390](#)
- [Date and Time Functions, page 391](#)

XPath Basics

XPath (XML Path Language) is an expression language developed by the World Wide Web Consortium (W3C) for addressing parts of XML documents. XPath also has basic manipulation functions for strings, numbers, and Boolean values.

To use XPath in TIBCO BusinessEvents, you need only be familiar with the basic XPath concepts, but you may wish to learn more about XPath when building complex expressions.

For a complete description of XPath, refer to the XPath specification (which can be obtained from www.w3.org).

TIBCO BusinessEvents uses XPath (XML Path Language) to identify elements whose content may be used, for, example in an event payload. You can also use XPath to perform basic manipulation and comparison of strings, numbers, and Boolean values.

Addressing Schema Elements

All Scope Variables and Function arguments are represented as an XML schema. Regardless of where the data comes from or its format, TIBCO BusinessEvents represents the data as a schema tree. The data can be simple (strings, numbers, Boolean values, and so on), or it can be a complex element. Complex elements are structures that contain other schema elements, either simple elements or other complex elements. Both simple and complex elements can also repeat. That is, they can be lists that store more than one element of the given type.

XPath is used to specify which schema element you would like to refer to. Each schema has its own associated structure, for example, a set of simple values or simple data and other complex data.

To reference a particular data item in a schema, you start with the root node and then use slashes (/) to indicate a path to the desired data element. For example, if you wish to specify the Street attribute in the ShipName complex element that is in the GetOrderInformation node, you would use the following syntax:

```
$GetOrderInformation/ShipName/Street
```

The path starts with a dollar sign to indicate it begins with a root node, then continues with node names using slashes, like a file or directory structure, until the desired location is named.

Evaluation Context

XPath also has a method for referencing relative paths from a particular node. If you have an *evaluation context*, or a particular starting node in a schema tree, you can specify the relative path to other elements in the tree.

For example, if your evaluation context is `$GetOrderInformation/ShipName`, then you can reference the sub-items of `ShipName` without specifying the entire path. If you wish to reference `$GetOrderInformation/RequiredDate`, the relative path would be `../RequiredDate`. The path is relative to the evaluation context — `RequiredDate` is one level higher in the schema tree than the elements of `ShipName`.

Namespaces

Some schema elements must be prefixed with their namespace. The namespace is automatically added to elements that require this when creating mappings or when dragging and dropping data in the XPath formula builder.

Search Predicates

An XPath expression can have a search predicate. The search predicate is used to locate a specific element of a repeating schema item. For example, a `$GetOrderInformation/OrderDetails/OrderItem` item is a repeating element. If you wish to select only the first item in the repeating element, you would specify the following:

```
$GetOrderInformation/OrderDetails/OrderItem[1]
```

The `[1]` specifies the first element of a repeating item.

Sub-items can also be examined and used in a search predicate. For example, to select the element whose `ProductId` is equal to "3A54", you would specify the following:

```
$GetOrderInformation/OrderDetails/OrderItem[ProductId="3A54"]
```

You can also use functions and expressions in the search predicate. For example, if you wish to find all elements after the first, you would specify the following:

```
$GetOrderInformation/OrderDetails/OrderItem[position()>1]
```

See the Functions tab of the XPath Formula Builder for a list of available functions available and online documentation.

Testing for Nil

Some elements can be explicitly set to nil. You can test an element to determine if it is set to nil or not. For example, the following XPath expression returns true if the `$Order/Item/OnSale` element is set to nil:

```
$Order/Item/OnSale/@xsi:nil="true"
```

Comments

You can add comments to XPath expressions using the XPath 2.0 syntax for comments. The syntax is:

```
{-- <comment here> --}
```

For example, the following XPath expression contains a comment:

```
$GetOrderInformation/ShipName/Street {-- returns the street --}
```

The XPath Formula Builder

You access the XPath Formula Builder using a button in the Function Argument Mapping Wizard. First select an item in the Function Argument panel (in the Input section). Then click the XPath Formula Builder button ()

The XPath formula builder allows you to drag and drop schema elements and XPath functions to create XPath expressions. The schema elements, when dragged into the XPath Formula field, automatically become valid XPath location paths for the desired item. If a function is dragged into the XPath formula window, there are placeholders for each parameter of the function. You can drag and drop schema elements over the parameter placeholders to replace each placeholder.

[Table 41](#) describes the different areas of the XPath formula builder.

Table 41 XPath Formula Builder Reference

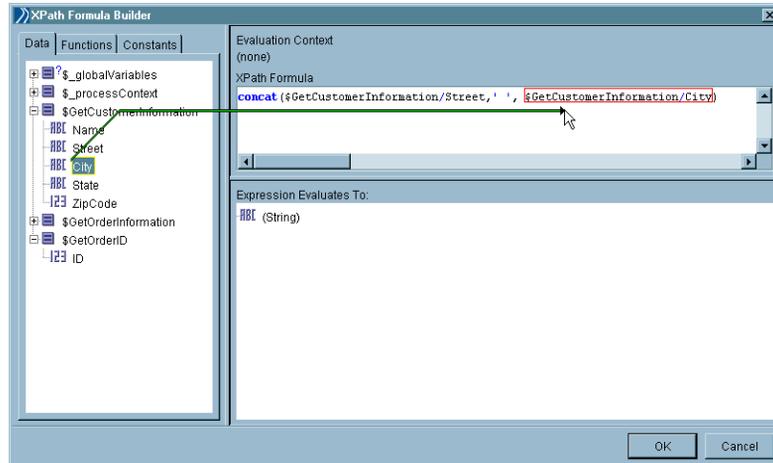
Element	Description
Data tab	Displays the Scope Variables schema tree. All elements in this tree are available to drag and drop into the XPath Formula field.
Functions tab	<p>Displays the available XPath functions. These are categorized into groups and each function can be dragged from the function list into the XPath Formula field.</p> <p>When the function is placed into the XPath formula, placeholders are displayed for the function's parameters. You can drag and drop schema elements from the Data tab into the function's placeholders.</p> <p>The result of evaluating the function is displayed in the "Expression Evaluates To" panel. If there are any errors in the expression, they are listed there as well.</p> <p>For more information about XPath functions, see the description of the function that is displayed when it is selected in the XPath formula builder.</p>

Table 41 XPath Formula Builder Reference

Element	Description
Constants tab	<p>Displays the constants available for use in XPath expressions. These are categorized into groups and each constant can be dragged from the constants list into the XPath Formula field.</p> <p>Constants are useful for inserting special characters, such as tabs, symbols, and so on, into XPath formulas. Constants are also defined for commonly used items, such as date formats.</p>
Documentation panel	<p>Appears below the Functions and Constants tabs. Describes each selected function. As you click on a function in the Function tab, the documentation panel gives a brief description of the function and one or more examples. Similarly documentation for constants in the Constants tab appears.</p>
Evaluation Context field	<p>Displays the evaluation context of the expression field that the editor was invoked from. See Evaluation Context on page 385 for more information about the evaluation context.</p>
XPath Formula field	<p>Displays the XPath formula you wish to create. You can drag and drop items from the Data tab or the Functions tab to create the formula.</p>
Expression Evaluates To Panel	<p>Displays the result of evaluating the formula shown in the XPath Formula field. If there are errors in the formula, they are displayed here.</p>

[Figure 6](#) illustrates using the XPath formula builder to create a valid function. The function concatenates the data elements `$GetCustomerInformation/Street` and `$GetCustomerInformation/City` and places a space between the two elements.

Figure 6 Creating an XPath formula



String Representations of Datatypes

When data must be represented in the input or output of an activity, the data is represented as a string. This section explains the string representations of various datatypes. TIBCO BusinessEvents follows the XPath 1.0 standard for representing all numeric datatypes. TIBCO BusinessEvents follows the XML Schema canonical format for all other datatypes.

Numeric Datatypes Numeric datatypes include all types derived from `xs:integer`, `xs:decimal`, `xs:float`, and `xs:double`.

All decimal, float, and double numbers are compressed to an integer when represented, if there are only zeros following the decimal point (for example, "1.000" is represented as 1). Scientific notation is never used to represent a floating point number as a string (for example, "`xs:double('1.234E05')`") is represented as 123400). Data is truncated if the number of digits exceeds the maximum precision for the datatype (for example, "`xs:float('1.23456789')`") is represented as 1.2345679).

Both zero and negative zero are represented as 0. Positive and negative infinity are represented as `Infinity` and `-Infinity`. Not a number is represented as `NaN`.

Boolean The boolean datatype is used to indicate a true or false state.

`xs:boolean('true')` and `xs:boolean('1')` are represented by `true`. The XPath function `true()` is also represented as `true`.

`xs:boolean('false')` and `xs:boolean('0')` are represented by `false`. The XPath function `false()` is also represented as `false`.

Date Datatypes TIBCO BusinessEvents Function Argument Wizard (also known as the function argument mapper) implements dates in one of two ways. Either a date is stored as the number of milliseconds since January 1, 1970, or the date is implemented according to the XPath 2.0 or XQuery 1.0 standards as a set of normalized components (`xs:date`, `xs:time`, `xs:dateTime`, and so on) with an optional time zone offset. Activities that are associated with Java (for example, Java Code, Java Method, and so on) use the first implementation. Activities that are associated with XML (for example, Mapper, Parse XML, and so on) use the second implementation. The second implementation supports arbitrary precision of the seconds component.

Conversion between these representations may result in a loss of information either because of the difference in time zone representation or the precision of the seconds.

Date and Time Functions

There are some functions in the XPath formula builder that allow you to parse or format strings that represent dates and times. These functions are:

- `format-dateTime(format, dateTime)`
- `format-date(format, date)`
- `format-time(format, time)`
- `parse-dateTime(format, string)`
- `parse-date(format, string)`
- `parse-time(format, string)`

The *format* parameter of these functions is based on the format patterns available for the `java.text.SimpleDateFormat` Java class. In the format parameter, unquoted alphabetic characters from A to Z and a to z represent the components of the date or time string. You can include non-pattern alphabetic characters in the string by quoting the text with single quotes. To include a single quote, use `'`. [Table 42](#) describes the alphabetic characters and their associated presentation in a date or time string.

Table 42 *Formatting characters in date or time strings*

Letter	Description	Example
G	Era Four or more Gs return the full name of the era.	AD
Y	year Two ys return two-digit year.	2003; 03
M	Month in year Three or more Ms return text name.	August; Aug; 08
w	Week in year	48
W	Week in month	3
D	Day in year	254
d	Day in month	28

Table 42 Formatting characters in date or time strings

Letter	Description	Example
F	Day of week in month	3
E	Day in week Four or more Es return the full name of the weekday.	Friday; Fri
a	AM/PM marker Four or more as return the full name.	AM
H	Hour in day (0-23)	23
k	Hour in day (1-24)	1
K	Hour in AM/PM (0-11)	11
h	Hour in AM/PM (1-12)	1
m	Minute in hour	59
s	Second in minute	48
S	Millisecond	456
z	Time zone represented as a GMT offset.	GMT-08:00
Z	RFC 822 four-digit time zone format	-0800
all other letters	Reserved	—

For any format pattern letter that returns a numeric value (for example, *w*, *h*, and *m*), the number of letters in the format pattern represents the minimum number of digits. For formatting functions, if the date or time has fewer digits than the number of pattern letters, the output is padded with zeros. For parsing functions, when the date or time has fewer digits than the number of characters in the format pattern, the extra characters are ignored, unless they are needed to determine the boundaries of adjacent fields.

Table 43 illustrates some example date and time format patterns and the resulting string.

Table 43 Example date and time format patterns

Date/Time Pattern	Result
"yyy.MM.dd G 'at' HH:mm:ss"	2003.3.11 AD at 09:43:56
"EEE, MMM d, ''yy"	Tue, Mar 11, '03
"hh 'o''clock' a, zzzz"	9 o'clock AM, GMT-8:00
"K:mm a"	0:08 PM
"yyMMdHHmmssZ"	010704120856-700

Chapter 23 **Cluster Deployment Descriptor (CDD)**

This short chapter provides general information about the Cluster Deployment Descriptor (CDD) and working with the CDD editor, which is used to configure CDD files.

Topics

- [Cluster Deployment Descriptor Overview, page 396](#)
- [How CDD Settings Apply at Runtime, page 400](#)
- [Setting Global Variables in the CDD File \(for Command Line Startup\), page 402](#)
- [Adding a Cluster Deployment Descriptor \(CDD\), page 403](#)
- [CDD Cluster Tab General Settings Reference, page 405](#)

Cluster Deployment Descriptor Overview



TIBCO BusinessEvents Express Content relating to Cache OM and backing store does not apply to TIBCO BusinessEvents Express edition.

The Cluster Deployment Descriptor (CDD) is an XML file used to configure a project for deployment. One EAR file and one CDD file define all the settings for all the engines and agents you want to deploy for a single application.

Because the deploy-time configuration settings for all processing units are in the CDD file, you do not have to rebuild the EAR file to make changes to deploy-time settings. However, as desired, you can use only the copy of the CDD file that is in the EAR, for tighter control and uniformity.

When you deploy a processing unit, you specify these items:

- An EAR file
- A CDD file that you have configured for that EAR
- A processing unit (engine) that is configured in the specified CDD file.

The CDD file you specify can be in the file system or in the EAR file. To specify a CDD located in an EAR file, provide its project path and name.



- For deployment using TIBCO Administrator, the CDD file in the EAR is used.
- The CDD file does not accept global variables as values.

You can configure multiple CDD files for a project for different purposes such as testing a design, trying out different object management options, dividing the work differently between agents and processing units (engines), and so on. However you use the same CDD file when deploying all the processing units for an application.

Maintaining the CDD File

The CDD is an XML file and you maintain it using the CDD editor.

Working with the CDD Editor

The CDD editor is organized into various tabs. Predefined fields make it easy to maintain common settings. Validation checks assist correct configuration. Settings for less common options are configured by adding properties to property sheets on the CDD editor's tabs.

Editing the CDD After Deployment

Edit the CDD file in Studio and copy the output to the deployment location as needed.

Great care must be taken to ensure that the deployed CDD is edited correctly and that all CDD files are kept in sync for straightforward project management. If an emergency change is made to the deploy-time file, ensure that such changes are added back to the design-time version of the file and maintained in TIBCO BusinessEvents Studio (and put under source control), so that subsequent redeployment can be done without the need for manual changes.

Direct Editing is Discouraged

Although the CDD file can be edited directly, doing so is strongly discouraged. The file is complex, and must adhere to the schema. The schema itself may change in a release, and migration is handled using TIBCO BusinessEvents Studio.

Defining and Configuring the Cluster and Object Management Type

Of the object management (OM) types available, the most commonly used type for production systems is Cache OM. Cache OM is generally used with a backing store for persistence of data generated in the engine. See *TIBCO BusinessEvents Architect's Guide* for an overview of clusters and object management types.

When you configure Cache OM, you must first choose a cache provider, TIBCO or Oracle. The internal TIBCO cache provider, TIBCO BusinessEvents DataGrid, is the default choice. Or you can use a supported version of Oracle Coherence, for which you have a license that is appropriate for your usage. See [Chapter 24, Cache OM and Cluster Configuration, on page 407](#).

For In Memory OM, no object management configuration is required because all objects are in memory and do not persist.

Configuring Management of Domain (Entity) Object Instances

If you choose Cache OM, you must also configure how to manage the domain (that is, entity) object instances. For example you can determine whether the instances are flushed from the Rete network after each RTC (as is generally recommended) or are kept in the cache.

If you set up a backing store, you can specify additional settings. For example you can define a subset of ontology object instances to be stored in the backing store, and you can control which (if any) object instances are loaded into the cache from the backing store at system startup.

See [Chapter 28, Domain Objects Configuration](#), on page 477 and [Chapter 26, Backing Store Configuration](#), on page 455.

Defining and Configuring Agent Classes and Processing Units

Cache OM uses multiple processing units, each running one or more agents. You define the agent classes and processing units in the CDD.

Agent Classes

Agent classes define the different sorts of agents you can deploy.

In Memory OM uses only inference agents, and each agent operates independently. With Cache OM, inference agents share the cache data.

For inference agent classes, you distribute a project's resources among the agent classes to define the specific work each agent will do.

The following types of agent classes require the use of Cache OM:

- Cache
- Query, used in TIBCO BusinessEvents Event Stream Processing
- Process, used in TIBCO BusinessEvents Process Orchestration
- Dashboard, used in TIBCO BusinessEvents Views
- Monitoring & Management, used by MM server only. (See Chapter 3, Basic MM Configuration of *TIBCO BusinessEvents Administration*.)

Collections Tab

At the Collections tab, you can (optionally) group rules, rule functions, and destinations into collections so that they can be easily assigned to agent classes, to simplify agent class configuration.

The Collections tab is also where you define log configurations, which are assigned to processing units (engines).



Startup and Shutdown Rule Functions At the Collections tab, put rule functions for use at start up (startup rule functions) into different groups from those used at shut down (shutdown rule functions) and arrange them in the order in which they are to be executed. Then you can select them appropriately at the agent classes tab.

See [Configuring Agent Classes \(All OM Types\)](#), page 496 and [Configuring Collections of Rules, Rule Functions, and Destinations](#), page 492.

Processing Units

Processing units deploy as engines within which the agents run. In the Processing Units tab, you define which agents to include in the processing unit, and which logging configuration to use. Depending on the OM, you also configure some additional settings.

See [Configuring Processing Units \(All OM Types\)](#), page 513.



For deployment using the TIBCO BusinessEvents Monitoring and Management component, processing units configured in the CDD editor are referenced in the Site Topology editor, where they are associated with physical resources. (See Chapter 3, Basic MM Configuration in *TIBCO BusinessEvents Administration*.)

How CDD Settings Apply at Runtime

It's important to understand the effect of design-time settings in the runtime environment. The tab and section within a tab where you set values in the CDD can affect the scope of those values, and how they can be overridden.

Using Properties at Different Levels

The scope of a property depends on the property sheet you add it to. Not all properties are valid at all levels. Use your judgment.

For example, properties that include the agent class name, such as `Agent.AgentClassName.checkDuplicates`, can be used at different levels. Here is the scope of each level for these `AgentClassName` properties:

Cluster level Applies to all `AgentClassName` agents in the cluster.

Processing unit level Applies to any `AgentClassName` agent deployed in the specified processing unit.

Agent class level Applies to any `AgentClassName` agent, used in any processing unit.

(Not all agent-level properties include the agent class in the property name.)

Only one value for a property is used when a processing unit is deployed. The order of precedence is explained in the section, [Order of Precedence at Runtime](#).

Order of Precedence at Runtime

This section also appears in TIBCO Administrator and is reproduced here for your information, because order of precedence at run time can affect decisions made at design time.

The order of precedence at run time, from highest priority to lowest, is as follows:

1. Command-line arguments at engine startup.
2. Properties set in property files specified at the command line.
3. Properties in the deployed TRA file.
4. CDD file, processing unit level (for the current PU).
 - a. Properties
 - b. Settings

5. CDD file, agent class level (for agents listed in the current PU settings, prioritized in reverse order of that list):
 - a. Properties
 - b. Settings
6. CDD file, cluster level:
 - a. Properties
 - b. Message encoding
 - c. Settings

EAR file properties (such as global variable overrides)



- Global variables set in the CDD file are ignored if you deploy using TIBCO Administrator. They are overridden by variables set in TIBCO BusinessEvents Monitoring and Management.
- Note that the TRA files should be used only for system-level settings that must be read before the JVM starts. All other properties should be in the CDD.

Setting Global Variables in the CDD File (for Command Line Startup)

Global variables are added to a project and their value is set using the TIBCO BusinessEvents Studio Global Variables editor, as explained in [Working with Global Variables on page 14](#).

You can also set global variable values in the CDD file. This can sometimes be useful if you will start the engine at the command line. It is not done for other purposes



Global variables set in TIBCO Administrator or TIBCO BusinessEvents Monitoring and Management (BEMM) override those set in the CDD If you will deploy using TIBCO Administrator or BEMM, set global variable values using the features provided in TIBCO Administrator or MM (whichever you will use), not in the CDD.

To Set Global Variables in the CDD

Add properties using this format:

```
tibco.clientVar.GVName = value
```

The *GVName* must exactly match the name set in the TIBCO BusinessEvents Studio Global Variables editor.



Global Variable Groups If global variables are defined in the TIBCO BusinessEvents project using groups, specify the group path using forward slashes. For example, if a variable *JMSuri* is located under a group called *URIs*, specify the variable as `tibco.clientVar.URIs/JMSuri`.

Add such properties at the appropriate level in the CDD, depending on the desired scope: cluster, processing unit, or agent class.

Adding a Cluster Deployment Descriptor (CDD)

When you add a CDD, you select an object management type (known as the OM type). The OM type determines many of the configuration options available in the CDD editor. For an introduction to this topic, see Chapter 6, Object Management Types in *TIBCO BusinessEvents Architect's Guide*.

Specifying and Changing the Object Management Type If you choose the OM type in the wizard, a template for that OM type is used. For example, if you choose Cache, then a cache agent with default values is created for you. If you choose the OM type in the editor (after you finish the wizard), you add the necessary settings yourself and defaults are not provided. You can switch back and forth between object types without losing any configuration values, but only until you save. For example, if you change from cache to memory and save, the cache agent and its configuration are lost.

For all OM types, however, the General settings are configured in the same way, as documented in this section.



Names in the CDD must conform to the NCName datatype. See the following page for more details:

<http://www.w3.org/TR/REC-xml-names/#NT-NCName>

Some Japanese characters, such as half-width Katakana, have issues when they are used in XML names. See the following document for more details:

<http://www.w3.org/Submission/japanese-xml/>

To Add a CDD

1. In BusinessEvents Studio Explorer, right click the folder where you want to store the CDD and select **New > Cluster Deployment Descriptor**. You see the New Cluster Configuration Wizard.
2. In the File name field, type a name for the CDD and click **Next**. (You can change the name in the editor as desired).



For Deployment using TIBCO Administrator By default, TIBCO Administrator looks for a processing unit named default and a CDD file named default.

3. At the Object Management Selection page, select an object management type for the deployment, then click **Finish**. You see the CDD editor.



Check the Object Management node on the left to be sure you selected the correct OM type. Depending on the OM type you chose, you see one of the following:

- **Object Management: [Cache]**
- **Object Management: [In Memory]**

You can right-click the Object Management element and select **Change to OM Type**.

4. Select Cluster tab > **General** on the left. On the right, specify the following:
 - The cluster name and message encoding.
 - As desired, an author name, and any comment you wish to record. (Version and date are not editable.)

See [CDD Cluster Tab General Settings Reference on page 405](#) for details.



For deployment with TIBCO Administrator The message encoding specified in the CDD file General settings must match the TIBCO Administrator domain's message encoding.

5. Save the resource.

The next steps for each object management type are provided in the following sections:

- In Memory OM: (No cluster configuration.) [Chapter 29, Agent and Processing Unit Configuration, on page 489](#).
- Cache OM: [Chapter 24, Cache OM and Cluster Configuration, on page 407](#)

Database Concepts Configuration

If you use database concepts, available in the TIBCO BusinessEvents Data Modeling add-on product, see *TIBCO BusinessEvents Data Modeling Developer's Guide* for details about configuring the Database Concepts node in the CDD.

Dashboard Agent Configuration

If you use the TIBCO BusinessEvents Views add-on, see *TIBCO BusinessEvents Views Developer's Guide* for details about configuring dashboard agents.

CDD Cluster Tab General Settings Reference

For the related procedure, see [Adding a Cluster Deployment Descriptor \(CDD\) on page 403](#).

Table 44 CDD Cluster Tab General Settings

Property	Notes
Cluster Name	<p>Specifies the name of the cache cluster.</p> <p>If TIBCO BusinessEvents DataGrid is the cluster provider, any spaces in the name are converted to underbar characters internally.</p> <p>Note Do not use the name <code>\$cluster</code>. It is a reserved name.</p> <p>Initially set to the CDD name.</p>
Message Encoding	<p>The encoding used in Rendezvous messages exchanged between TIBCO applications.</p> <p>For deployment with TIBCO Administrator The message encoding specified here must match the TIBCO Administrator domain's message encoding.</p> <p>Default is ISO8859-1 (which is also the default for TIBCO Administrator)</p>
Author	<p>The name of the author of this CDD, as desired.</p> <p>Initially set to the currently logged-on user name.</p>
Comment	<p>Any comments as desired. Comments persist across versions.</p>
Version	<p>View-only field to record the version of the CDD, for information only. You could, for example, check whether deployed CDDs are all using the same version.</p>
Date	<p>View-only field to record creation time of this version.</p>



TIBCO BusinessEvents Express This chapter relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

Use this chapter if you picked Cache OM in the second page of the New Cluster Configuration wizard. This chapter explains how to configure cluster discovery and internal communications, and how to configure cache-related object management settings.

Topics

- [Cache Object Management Configuration, page 408](#)
- [CDD Cluster Tab and Cache OM Settings Reference, page 411](#)
- [Synchronous and Asynchronous Replication of Cache Objects, page 414](#)
- [Property for Cache Based Object Management on AIX, page 415](#)
- [Configuring Cluster Discovery and Internal Communication, page 416](#)
- [Configuring DataGrid Discover URL, page 419](#)
- [Configuring DataGrid Listen URL, page 422](#)
- [CDD Cluster Tab DataGrid Properties Reference, page 430](#)
- [Enabling Use of Oracle Coherence as the Cache Provider, page 433](#)
- [Configuring Oracle Coherence Cluster Discovery, page 434](#)
- [CDD Cluster Tab Coherence Properties Reference, page 438](#)

Cache Object Management Configuration

When you configure object management for Cache OM, you select the cache provider, configure a cache agent quorum, number of backup copies. If you have a backing store, you can also configure a limited cache and specify the cache size for each entity.

In addition to the basic reference tables (as noted in the procedure below) the following sections have additional specific guidelines:

- [Configuring a Limited \(or Unlimited\) Cache on page 408](#)
- [Synchronous and Asynchronous Replication of Cache Objects on page 414](#)
- [Property for Cache Based Object Management on AIX on page 415](#)
- [Configuring Cluster Discovery and Internal Communication on page 416](#) and sections following

For backing store details, see [Chapter 26, Backing Store Configuration, on page 455](#) for project configuration details and see [Chapter 30, JDBC Backing Store Setup, on page 523](#) for details about setting up the backing store itself.

To Configure Cache OM Options

1. Add a CDD file or open the CDD file you added, as explained in [Adding a Cluster Deployment Descriptor \(CDD\) on page 403](#).
2. Select the Cluster tab > Object Management node on the left and on the right, configure settings as explained in [CDD Cluster Tab and Cache OM Settings Reference on page 411](#).
3. Save the resource.

Configuring a Limited (or Unlimited) Cache

This procedure explains how to use the options available to configure a cache of a limited size.

You can also use this procedure to set options for an unlimited cache.

Also see Limited and Unlimited Cache Size in *TIBCO BusinessEvents Architect's Guide*.



Use of limited cache is supported only when a backing store is also used. The backing store retains entries in excess of the limit. Without use of a backing store data inconsistencies could result:

- Entries for an object in the object table (an internally used cache) and in the object cache itself could expire independently of each other.
- Domain object settings for limited cache apply at the object level. Related concepts could have different settings. For example, a container concept could have a limited cache setting and its container concept an unlimited cache setting. Each could be evicted at different times.

As desired, you can set the cache to limited at the default level and unlimited for specified objects; or you can set the cache to unlimited at the default level and limited for specified objects.



When the cache is limited, the number of cache objects is $\text{limit} * \text{no. of cache servers}$.

To Configure Limited Cache and Object Table Cache Options

1. Select the CDD editor Cluster tab > Object Management node.
2. In the Entity Cache Size setting, enter the desired number of objects per entity type.
3. In the Object Table Cache Size, enter the desired number of objects (handles) in the object table cache. You cannot set this value differently for different object types. See *The Role of the Object Table*, in *TIBCO BusinessEvents Architect's Guide*, for more details about the object table.
4. Select the Cluster tab > Domain Objects > Defaults node. Check the Is Cache Limited checkbox to enable limited cache globally. (Or uncheck the checkbox to use an unlimited cache globally.)
5. Select the Domain Objects > Overrides node. Select an override entry (or add one as needed).
6. Set the Is Cache Limited checkbox for the selected object type in one of the following ways:
 - If limited cache is set at the default level, uncheck the overrides Is Cache Limited checkbox to use an unlimited cache for objects of this type.
 - If unlimited cache is set at the default level, check the Is Cache Limited checkbox to use a limited cache for objects of this type.

See [Configuring a Limited \(or Unlimited\) Cache on page 408](#) for important requirements. For reference documentation on the Is Cache Limited checkbox see [Table 53, CDD Cluster Tab Domain Object Default Settings, on page 481](#).

Ensuring Multiple Clusters Do Not Conflict

With TIBCO BusinessEvents DataGrid clusters, use a different value for Cluster Name (in the Cluster tab, General node) and also use different discovery values.

With Coherence clusters, use a different value for cluster name.

CDD Cluster Tab and Cache OM Settings Reference



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

For the related procedure, see [Configuring Cluster Discovery and Internal Communication on page 416](#).

Table 45 CDD Cluster Tab Cache OM Settings (Sheet 1 of 3)

Property	Global Variables	Notes
Provider	No	<p>Select the cache provider:</p> <p>TIBCO: Uses the provided TIBCO BusinessEvents DataGrid component. See Configuring DataGrid Discover URL on page 419 and Configuring DataGrid Listen URL on page 422 for more details.</p> <p>ORACLE: Uses a supported version of Oracle Coherence, for which you have a license that is appropriate for your usage. Requires additional configuration. See Enabling Use of Oracle Coherence as the Cache Provider on page 433 and Configuring Oracle Coherence Cluster Discovery on page 434 for details.</p>
Cache Agent Quorum	Yes	<p>Specifies a minimum number (quorum) of storage-enabled nodes that must be active in the cluster when the system starts up before the following occur:</p> <ul style="list-style-type: none"> • Data is preloaded from the backing store, if a backing store is configured and preloading is configured (see Chapter 26, Backing Store Configuration, on page 455). • The other agents in the cluster become fully active. <p>The property does not affect the running of the deployed application after startup (though a message is written to the log file if the number of cache agents running falls below the number specified in this property). As a guideline, set to the number of cache agents configured.</p> <p>Default is 1.</p>

Table 45 CDD Cluster Tab Cache OM Settings (Sheet 2 of 3)

Property	Global Variables	Notes
Number of Backup Copies	Yes	<p>The number of backup copies (also known as the backup count) specifies the number of members of the distributed cache service that hold the backup data for each unit of storage in the cache. Recommended values are 0, 1, or 2.</p> <p>Value of 0 means that in the case of abnormal termination, some portion of the data in the cache will be lost. Value of N means that if up to N cluster nodes terminate at once, the cache data will be preserved. A backup count of 1 means one server plus one backup are needed, that is, two cache agents (or storage enabled nodes if cache agents are not used).</p> <p>To maintain the partitioned cache of size M, the total memory usage in the cluster does not depend on the number of cluster nodes and will be in the order of $M*(N+1)$.</p> <p>See Synchronous and Asynchronous Replication of Cache Objects on page 414 for details on replication behavior and options for TIBCO BusinessEvents DataGrid.</p> <p>Default is 1.</p>
Entity Cache Size	Yes	<p>Specifies the size of the limited cache, in number of cache entries for each object type. The setting is per processing unit. See Configuring a Limited (or Unlimited) Cache on page 408.</p> <p>Default is 10000 (entries per object type)</p>
Object Table Cache Size	Yes	<p>Specifies the maximum size of the object table cache, in number of entries. Used with limited cache only.</p> <p>See <i>The Role of the Object Table in TIBCO BusinessEvents Architect's Guide</i> for more details about the object table.</p> <p>Also see Configuring a Limited (or Unlimited) Cache on page 408.</p> <p>Default is 100000 entries</p>

TIBCO BusinessEvents DataGrid

Discovery URL	Yes	Specifies how an engine (node) listens for discovery requests from nodes attempting to join the cluster. PGM protocol is supposed for multicast discovery. TCP Protocol is supported for unicast discovery
---------------	-----	--

Table 45 CDD Cluster Tab Cache OM Settings (Sheet 3 of 3)

Property	Global Variables	Notes
ListenUrl	Yes	The discovery mechanism is based on pure TCP. All the designated well known metaspace members are identified by an IP address and a port number.
Remote Listen URL	Yes	Specifies on which IP address and TCP port this proxy metaspace member will be listening for the remote client connections.
Protocol Timeout	Yes	Indicates the protocol timeout value for space. The protocol can be unicast or multicast. Default value is -1 (forever).
Read Timeout	Yes	Indicates the read timeout value for the space, if a read timeout has been set. It specifies the read timeout value for a specified SpaceDef. The read timeout value applies to Get operations. Default value is 60000 (ms).
Write Timeout	Yes	Indicates the write timeout value for the space, if a write timeout has been set. Specifies the write timeout value that is set for the space. The write timeout value applies to Put, Take, Lock, and Unlock operations. Default value is 60000 (ms).
Lock Timeout	Yes	For a space that is locked, it specifies how long a member process will wait for it to become unlocked. The default is -1 (forever). Other valid values are 0 or any positive value.
Shutdown wait	Yes	Indicates the shutdown wait value for the space. Default is 8500 (ms).
Worker Thread Count	Yes	Indicates the thread count specified for the space. Default value is 4.
Store Properties as Individual Fields	No	Enabling this property sets the property <code>be.engine.cluster.as.tuple.explicit</code> to true.

Synchronous and Asynchronous Replication of Cache Objects

When you set a backup count, you define the number of backup object copies to make, in addition to the primary cache object. Backup cache writes can be done synchronously or asynchronously. There is a difference between the behavior of TIBCO BusinessEvents DataGrid and for Oracle Coherence as regards replication.

TIBCO BusinessEvents DataGrid Uses Asynchronous Replication

TIBCO BusinessEvents DataGrid is set up to use asynchronous replication. There is no option to use synchronous replication. Asynchronous replication allows you to run tests using a single cache agent.

With asynchronous replication, the inference agent writes to a cache agent and returns. The cache provider then makes a separate call to another cache agent to make the replica. This means that the writes from the inference agent do not incur the cost of synchronous replication, because replication happens on a different thread in the background. However, a small window exists in which the inference agent has written to the cache, and the cache provider has not replicated the data yet. If the cache agent fails at this point, data is lost because there is no replica. To safeguard the data, use a backing store with cache-aside database write strategy.

Oracle Coherence Uses Synchronous Replication

With Oracle Coherence as the cache provider, TIBCO BusinessEvents is set up to use synchronous replication. There is no option to use asynchronous replication.

With synchronous replication, when the inference agent writes to the cache, the cache provider makes a network call to another cache agent and makes a replica (or replicas) and then the call from the inference agent returns. This means making two cache puts for each cache operation. Therefore, synchronous replication is slower than asynchronous replication.

You must also ensure that the required number of cache agents is always up and running (depending on the backup count).

Property for Cache Based Object Management on AIX

When TIBCO BusinessEvents is installed on AIX and uses cache-based object management, you must add this property to all TRA files, and set the value to true.

```
java.net.preferIPv4Stack=true
```



Remember to set this property on all internal TIBCO BusinessEvents engines' TRA files too, such as in `be-mm.tra` for the TIBCO BusinessEvents Monitoring and Management (MM) server and the MM broker properties set in the MM CDD file. Add-on products also have engine TRA files you must update.

If you do not add this property, you see the following exception:

```
java.net.SocketException: The socket name is not available on this system
```

Configuring Cluster Discovery and Internal Communication

When you add a CDD file and select Cache OM type, you must configure how the members of the cache cluster discover each other at runtime and communicate with each other once the cluster is established.

This section has summary steps for both cache providers, and pointers to sections with more details.

To Configure a TIBCO BusinessEvents DataGrid Cluster (Metaspace)



An active LAN connection (device enabled and network cable plugged in) is required for TIBCO BusinessEvents DataGrid to work.

1. Add a CDD file or open the CDD file you added, as explained in [Adding a Cluster Deployment Descriptor \(CDD\) on page 403](#).

2. Select the Cluster tab > Properties node on the left and on the right, add the following two properties, as needed:

```
be.engine.cluster.as.discover.url
be.engine.cluster.as.listen.url
```

The properties can be omitted if you use PGM multicast with default values. See the following sections for details on configuring these properties:

- [Configuring DataGrid Discover URL on page 419](#)
- [Unicast \(Well-Known Address\) Cluster Member Discovery on page 420](#)
- [Configuring DataGrid Listen URL on page 422](#)

3. If you use unicast (well-known address) discovery, and you use TIBCO BusinessEvents Monitoring and Management for monitoring and management, you must also do the following (in the to-be-monitored project CDD):

- a. Add the following property to the cluster properties sheet.

```
be.mm.cluster.as.listen.url MMHostIP: Port
```

Specify the IP of the computer hosting the MM server, and an unused port.

- b. Add the value of the `be.mm.cluster.as.listen.url` property to the list of addresses in the `be.engine.cluster.as.discover.url` property, which should be present at the cluster level (so the value is identical for all potential cluster members).

The discover URL for well-known address configuration uses the following format:

```
tcp://ip:port[;ip:port]*
```

To Configure an Oracle Coherence Cluster

1. Before you begin ensure that you have met all prerequisite steps. See [Enabling Use of Oracle Coherence as the Cache Provider on page 433](#)
2. Add a CDD file or open the CDD file you added, as explained in [Adding a Cluster Deployment Descriptor \(CDD\) on page 403](#).
3. Select the Cluster tab > Properties node on the left and on the right, add properties as explained in [Configuring Oracle Coherence Cluster Discovery on page 434](#).
4. Add any other cluster level properties as needed. See [Other Coherence Properties on page 442](#).

To Configure Cache-Related Object Management Settings

1. Open the CDD file you added, as explained in [Adding a Cluster Deployment Descriptor \(CDD\) on page 403](#).
2. Select the Cluster tab > Object Management node on the left, and on the right configure settings as explained in [CDD Cluster Tab and Cache OM Settings Reference on page 411](#). Also see [Synchronous and Asynchronous Replication of Cache Objects on page 414](#).

Support for Hostaware Replication

Hostaware replication for ActiveSpaces is controlled by the property `be.engine.cluster.as.hostaware.enable`.

Hostaware replication requires that the member name be a 2-part name separated by a "."

By default, this property is `true` (or enabled).

- If `true`, the ActiveSpaces member name will be set as:

```
hostname.be-engine-name
```

where `be-engine name` is what is given on the `-n` command line option

When hostaware replication is enabled, if the cache nodes are not deployed on multiple machines to satisfy replication by the host, then replication will not happen (or will happen only according to the number of hosts available).

For example, if "Number of Backup Copies" is set to "1" and all cache nodes are deployed on a single host, then replication will be disabled (regardless of the number of cache nodes on that single host). If "Number of Backup Copies" is set to "2", and cache nodes are deployed only on 2 hosts, then only "1" backup copies will be maintained.

- If the property is `false`, hostaware replication will be disabled and the ActiveSpaces member name will be set as *be-engine-name*.

Disabling hostaware replication will honor "Number of Backup Copies", provided that there are enough cache nodes deployed in the cluster.

Configuring DataGrid Discover URL

When a cluster starts up, and also when new members join a cluster, a discovery process enables the members to discover each other. The discover URL specifies how an engine (node) listens for discovery requests from nodes attempting to join the cluster.

After the discovery is complete, the members communicate internally using a listen URL (explained in [Configuring DataGrid Listen URL on page 422](#)).

Two types of discovery are available:

- Multicast discovery (PGM)
- Unicast discovery (TCP), also known as "well-known address" discovery

Configuration for both discovery methods is explained below.



A TIBCO BusinessEvents DataGrid cluster is also known as a *metaspace*.

A TIBCO BusinessEvents engine is a *node* in the metaspace.

If No Other Cluster Members are Started

If a newly started node does not discover any running cluster nodes, the behavior is different depending on the type of discovery used:

- If multicast discovery is used, the newly started node becomes the first node of a newly started cluster.
- If unicast (well-known-address) discovery is used there are two cases:
 - If the address of the newly started node is not in the discover URL's list then it continues to wait for other well-known nodes to start, and a warning is written to the console while it waits.
 - If the address of the newly started node is in the discover URL's list, then it becomes the first node of a newly started cluster.

Multicast (PGM) Cluster Member Discovery

The discover URL for multicast discovery uses PGM (Pragmatic General Multicast) protocol.

The discovery property is `be.engine.cluster.as.discover.url`. For multicast discovery, the value is a URL with the following format:

```
tibpgm://destinationPort/network/
```

The default values equate to the following: `//7888/;239.8.8.9/`
Specify the parameters as follows.

Parameter	Notes
<i>destinationPort</i>	<p>Specifies the destination port used by the PGM transport.</p> <p>Must be the same value on all machines in the cluster.</p> <p>Default value is 7888.</p>
<i>network</i>	<p>Specifies the IP address of the interface to be used for sending multicast packets, and the multicast group address to be used.</p> <p>The format is as follows: <i>interface;multicast group address</i></p> <p>The value for <i>interface</i> is unique to a node. It must also be the same in both the discovery and the listen URLs for a node. If there are multiple interfaces on one machine, specify the interface you want to use and do not rely on the default value.</p> <p>The value for <i>multicast group address</i> must be the same on all machines in the cluster.</p> <p>The default value for <i>interface</i> is the first available interface provided by the operating system hosts file for the machine.</p> <p>Note If the desired interface is not listed in the hosts file then PGM picks the first available interface in the file. (On most operating systems, this file is called the <code>/etc/hosts</code> file.) If the first interface is the loopback interface (127.0.0.1) then PGM fails to start. In this case you would see a stacktrace exception in the log file such as the following:</p> <pre>SYS_ERROR (multicast_error - (8) grp_iface not a valid multicast interface)</pre> <p>To resolve this issue, either modify the hosts file, or provide the desired interface explicitly in the <i>network</i> argument.</p> <p>The default value for <i>multicast group address</i> is the multicast group address <code>239.8.8.9</code>.</p>

Unicast (Well-Known Address) Cluster Member Discovery

If you cannot or do not wish to use multicast discovery in your environment, then configure unicast discovery, also known as "well-known address" or WKA discovery. These "well-known addresses" enable a newly started node to discover existing members. Unicast discovery uses the TCP protocol.

The discovery property is `be.engine.cluster.as.discover.url`. For unicast discovery, the value is a semicolon-separated list comprising a sub-set of all the listen URLs (which are different for each PU), using this format:

```
tcp://ip:port[;ip:port]*/
```



One cluster node in the WKA list must be running at all times At least one cluster node specified in the well-known address list must be running at all times, so that other new members can join the cluster (metaspace). If all nodes specified in the well-known address list stop, then other nodes that are still running continue to function, but they print warnings to the console and no new members can connect to this cluster.

For WKA discovery, make discover URL a cluster-level property and listen URL a PU-level property The discover URL property

(`be.engine.cluster.as.discover.url`) must be present and configured identically for all potential cluster members. Therefore add this property at the cluster level of the CDD file. The listen URL property

(`be.engine.cluster.as.listen.url`) must be present and configured differently for each possible cluster member. Therefore add this property at the PU level.

Configuring DataGrid Listen URL

The listen URL is used for direct communication between the members of the metaspace. It is configured the same way for multicast and for unicast discovery (see [Configuring DataGrid Discover URL on page 419](#)). The listen URL value must be different for each cluster member, so configure it at the PU level.

The listen URL uses this format:

```
tcp://interface:port[-EndPort |*]/
```

The cluster member binds to the specified interface and the specified port when creating the TCP socket. Specify the parameters as follows.

Parameter	Notes
<i>interface</i>	<p>To specify a value, use the desired IP address.</p> <p>The value for <i>interface</i> must be the same in both the discovery and the listen URLs for a node. If there are multiple interfaces on one machine, specify the interface you want to use and do not rely on the default value.</p> <p>The default value for <i>interface</i> is the first available interface provided by the operating system for the machine.</p>
<i>port</i>	<p>To specify a single port use the port number in the listen URL, as shown in this example:</p> <pre>tcp://interface:6000/</pre> <p>You can use an auto-incrementing feature, as explained in Auto-incrementing Within a Range of Ports, page 423.</p> <p>The default value is the first available port in the 50000+ range.</p>

Multiple Nodes on One Machine

If multiple nodes (engines) are running on one machine, identify each uniquely. Use the same value for *interface*, but a different value for *port* for each node.

Auto-incrementing Within a Range of Ports

If a machine has blocked some ports in the default range, or if you want to use a different range, you can configure the listen URL to start with a specified IP address and port, and optionally provide an upper limit. If the specified port is not available, TIBCO BusinessEvents auto-increments the port until it finds an available port, up to the specified upper limit, if any. To specify a specific range use this format:

```
tcp://interface:Port-EndPort/
```

For example, given the following listen URL, TIBCO BusinessEvents attempts to open port 8000 and if it is not available it tries the next port number, until it finds an available port, up to 9000 (inclusive). If none is available, it keeps retrying. Make some ports in the specified range available so that the cluster nodes can start.

```
tcp://interface:8000-9000/
```

To specify a range with the upper limit of unsigned short minus one, use this format:

```
tcp://interface:Port-*/
```

Configuring DataGrid Transport Security

Transport-level security allows you to protect data being transported within the DataGrid by preventing these issues:

- Alteration of traffic
- Eavesdropping
- Exchange of data between untrusted parties

When security is used, any transmission of messages within the DataGrid occurs on a secure transport. A security domain's `transport_security` setting controls the level of security used for communication within the DataGrid.

The available settings for `transport_security` are:

- **encrypted_normal** Use secure transport with 128 bit symmetric key encryption (default).
- **encrypted_strong** Use secure transport with 256 bit symmetric key encryption.
- **integrity** Use secure transport without encryption.

The two possible node types in a secure DataGrid are:

- **Controllers:** Nodes dedicated to enforcing a security domain's defined security behavior for a cluster associated with the security domain. Security domain controllers are the only discovery nodes in a cluster.
- **Requestors** Nodes that require access to the data in the DataGrid, such as a seeder or a leech, and which need to be authorized by a controller. Requestors can never be used a discovery nodes.

Set Up Policy and Token Files

Follow these steps to set up policy and token files for a given cluster:

1. Start BusinessEvents with a non-secure DataGrid cluster.
2. Create a security policy file using the ActiveSpaces admin utility (`TIBCO_HOME/as/2.1/bin`).

```
as-admin> create security_policy
policy_name "mypolicy/mydomain"
policy_file "mypolicy.txt"
```

3. Edit the Metaspaces Access List for the security domain in the security policy file.
 - Set the metaspaces name (no quotes)
 - Set the discovery URL of the metaspaces which can be TCP discovery only (no quotes)
4. Ensure that there is a `metaspaces_access` entry with the cluster name from the BusinessEvents CDD.
5. Review the Transport Security settings to ensure they are set to meet your security requirements.
6. Save your changes to the security policy file.
7. Validate your security policy file using the ActiveSpaces admin utility.


```
as-admin> validate policy_name "mypolicy" policy_file
"mypolicy.txt"
```
8. Create a security token file using ActiveSpaces admin utility


```
as-admin> create security_token
domain_name "mydomain"
policy_file "mypolicy.txt"
token_file "mytoken.txt"
```
9. Ensure that there is a `metaspaces_access` entry with the cluster name from the BusinessEvents CDD.
10. Validate your security token file using the ActiveSpaces admin utility:


```
as-admin> validate token_file "mytoken.txt"
```
11. Shut down the cluster.
12. Enable security by make the following property true in the CDD:


```
be.engine.cluster.as.security.enable=true
```
13. By default, TIBCO BusinessEvents assumes all nodes to be requestors. However, every cluster must have at least one controller node.

You can change this by setting the role property in the CDD:

```
be.engine.cluster.as.security.mode.role=Requester or Controller
(the value is case-insensitive)
```
14. Supply the security file with the following property in the CDD:


```
be.engine.cluster.as.security.file=<File location of Controller
(policy file) or Requester (token file)>
```

EXAMPLE:

Suppose you start two nodes in a cluster called `mycluster`. One is inference engine and the other is a cache engine.

1. Assume that the `LustenURL` for the cache (controller) is `tcp://10.98.192.101:9091` and the `LustenURL` for the inference (requester) is `tcp://10.98.192.101:9090`.
2. Generate the policy and token files from `as_admin`. Edit the `metaspace_access` line in both files to set the correct cluster name.
3. Also edit the line to place the listen url of the controller into the `discovery=` setting.

```
metaspace_access=metaspace=mycluster;discovery=tcp://10.98.192.101:9091;
```
4. For the inference engine, you can then set the following properties:

```
be.engine.cluster.as.security.enable=true
be.engine.cluster.as.security.mode.role=Requester
be.engine.cluster.as.security.file=C:/temp/mytoken.txt
```
5. For the cache engine, you can then set:

```
be.engine.cluster.as.security.enable=true
be.engine.cluster.as.security.mode.role=Controller
be.engine.cluster.as.security.file=C:/temp/mypolicy.txt
```



There is no rule that caches have to be controllers, or inferences have to be requesters. The roles are interchangeable as long as there is one controller in the cluster.

Restricting Transport Access

Transport level security allows you to restrict transport connections within a security domain to only the “trusted” nodes.

To restrict transport connections within a security domain:

1. Open the security policy file for the domain in a text editor
2. Go to the line that reads `transport_access=false;cert_file=`
3. Edit the line to read:

```
transport_access=true;cert_file=<trusted_certs_file>
where trusted_certs_file is the filename for a trusted
certificate file that you will create in step 8.
```

4. Save the security policy file.

5. Use the `validate policy_file` command to validate the security policy file.
6. Use the ActiveSpaces admin utility to generate a security token file from the security policy file, which contains its own private key and a public certificate. This key and certificate are used to verify the identity of a node using the security token file when it tries to initiate any transport connections. For example:


```
as-admin> create security_token
domain_name "mydomain"
policy_file "mypolicy.txt"
create_identity
token_file "mytoken.txt"
```
7. Use the `validate token_file` command to validate the security token file.
8. Create an empty trusted certificates file to hold the public certificates of the nodes to allow transport connections from.
9. Copy and paste the public certificate of the local token identity from the security token file into the trusted certificates file.

The public certificate is everything in the security token file between and including

```
-----BEGIN CERTIFICATE-----
-----END CERTIFICATE-----
```

10. Save the trusted certificates file.
11. Start a security domain controller using the security policy file name when connecting to the DataGrid

DataGrid communication within the security domain is now restricted to only security domain controllers and security domain requestors that connect to the cluster using a security token file whose public certificate is contained in the trusted certificates file.

Schema/Model Migration with Shared Nothing Persistence

You can choose one of these options:

1. Deploy the new EAR with the additional field and restart all nodes. System will automatically alter the space as needed during recovery.
2. If it is not desired to restart cache nodes, shut down all inference engines. Connect with AS administrator and alter the space to add new fields as shown in example below.

The new field must be nullable as is the case for all user fields in BusinessEvents .

```
alter space
name"dist-unlimited-bs-readOnly-Test--be_gen_Concepts_Simple" add
(field name "long_field" type 'LONG' nullable true)
```

```
alter space
name"dist-unlimited-bs-readOnly-Test--be_gen_Concepts_Simple" add
(field name "con_cept_array" type 'BLOB' nullable true).
```

Then start the inference engines.

3. Hot deployment of new properties for Shared Nothing persistence has been added .



Adding a concept property of type contained concept with the contained concept type set to an existing concept is supported only for option 1. and 2.

4. Dropping existing field

This option is available only in AS-2.1.2 release. Keep the existing unused fields as is, without assigning any values.



With AS-2.1.2: simply deploy the new EAR without the dropped field and restart all the nodes. System will automatically alter the space as needed during recovery.

5. Altering existing field types: Not supported (requires a custom solution).

Hot Deployment of New Properties

Hot deployment of new properties into existing TIBCO BusinessEvents concepts is available only in these cases:

- When cache Object Management is enabled with no persistence or Shared Nothing persistence.

- When the concepts with new properties are cache-only.

To Enable hot deployment of new properties:

Before starting TIBCO BusinessEvents Studio, do the following:

1. Add the following line to the vmargs in the `studio.ini` file:
`-Dcom.tibco.be.hotdeploy.concept=true`
2. Set the property `com.tibco.be.hotdeploy.concept=true` in the CDD file.

Note the following restrictions in setting this functionality::

- When enabling it with no backing store, ensure that the property `be.engine.cluster.as.tuple.explicit=true` is set in the CDD file.
- This property is set to `true` by default when using Shared Nothing persistence.
- Adding a concept property of type `contained concept` with the `contained concept` type set to an existing concept is not supported.

CDD Cluster Tab DataGrid Properties Reference



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.



Discovery and listen URL interfaces must match Ensure that a node's interface (IP address) is specified using the same value in the discover URL and in the listen URL. If there are multiple interfaces on one machine specify the IP explicitly in both properties.

Table 46 CDD Cluster Tab TIBCO BusinessEvents DataGrid Properties

Property	Notes
<code>be.engine.cluster.as.discover.url</code>	<p>The discover URL specifies how an engine (node) listens for discovery requests from nodes attempting to join the cluster. PGM protocol is supported for multicast discovery. TCP protocol is supported for unicast (well-known address) discovery. Configuration is different for multicast and unicast discovery. See Configuring DataGrid Discover URL on page 419 for details.</p> <p>The default value for multicast equates to: <code>tibpgm://7888/;239.8.8.9/</code></p>
<code>be.engine.cluster.as.hostaware.enable</code>	<p>By default, this property is true (or enabled).</p> <ul style="list-style-type: none"> If true, the ActiveSpaces member name will be set as: <code>hostname.be-engine-name</code> where be-engine name is what is given on the <code>-n</code> command line option <p>When hostaware replication is enabled, if the cache nodes are not deployed on multiple machines to satisfy replication by the host, then replication will not happen (or will happen only according to the number of hosts available).</p> <p>For example, if "Number of Backup Copies" is set to "1" and all cache nodes are deployed on a single host, then replication will be disabled (regardless of the number of cache nodes on that single host). If "Number of Backup Copies" is set to "2", and cache nodes are deployed only on 2 hosts, then only "1" backup copies will be maintained.</p> <ul style="list-style-type: none"> If the property is false, hostaware replication will be disabled and the ActiveSpaces member name will be set as be-engine-name. <p>Disabling hostaware replication will honor "Number of Backup Copies", provided that there are enough cache nodes deployed in the cluster.</p>

Table 46 CDD Cluster Tab TIBCO BusinessEvents DataGrid Properties (Cont'd)

Property	Notes
<code>be.engine.cluster.as.listen.url</code>	<p>The listen URL is used for direct communication between the members of the metaspace after the discovery process. The listen URL uses this format:</p> <pre>tcp://interface:port/</pre> <p>You can also use an auto-incrementing feature by specifying a range as follows:</p> <pre>tcp://interface:Port-[toPort *]/</pre> <p>The default value for <i>interface</i> is the first available interface provided by the operating system for the machine. See Configuring DataGrid Listen URL on page 422 for details.</p> <p>The default value for port is the first available port in the 50000+ range.</p>
<code>be.engine.cluster.as.security.enable</code>	<p>Enable Transport level security for the DataGrid.</p>
<code>be.engine.cluster.as.security.file</code>	<p>Full path to the policy file in case the node is a Controller or the token file in case the node is a Requester</p>
<code>be.engine.cluster.as.security.mode.role</code>	<p>Security role of a node for the secure DataGrid. Possible values are: <i>Controller</i> or <i>Requestor</i>.</p> <p>The Controller is dedicated to enforcing security behavior for a cluster associated with the security domain. Security Controllers are the only discovery nodes in a cluster.</p> <p>The Requestor requires access to the data in the DataGrid, which needs to be authorized by the Controller. A Requestor can never be used as a discovery node.</p>

Table 46 CDD Cluster Tab TIBCO BusinessEvents DataGrid Properties (Cont'd)

Property	Notes
<code>be.mm.cluster.as.listen.url</code>	<p data-bbox="315 291 1276 387">If you use well-known-address discovery and TIBCO BusinessEvents Monitoring and Management (MM), you must also add this property at the cluster level of the to-be-monitored project's CDD file.</p> <p data-bbox="315 404 1276 465">Specify the value as the IP of the computer hosting the MM server, and an unused port.</p> <p data-bbox="315 482 1276 543">See Configuring Cluster Discovery and Internal Communication on page 416 for the procedure, including an additional step you must take.</p> <hr data-bbox="97 560 1276 569"/>

Enabling Use of Oracle Coherence as the Cache Provider

If you want to use the Oracle Coherence cache provider, you must provide a fully licensed, supported version of the software. See the product readme file for supported version information. You must also enable Oracle Coherence as the cache provider, as explained below.



You cannot use the `coherence.jar` file from an earlier TIBCO BusinessEvents release.

To Enable Oracle Coherence as the Cache Provider

1. Copy the `coherence.jar` file from your Oracle Coherence installation to `BE_HOME/lib/ext/tcpl`.

(This location is preconfigured in the `studio/eclipse/configuration/studio.tra` classpath, as shipped. If you use a different location, update the classpath.)
2. Enable the Coherence category functions as follows:
 - a. Open the `BE_HOME/studio/eclipse/configuration/studio.tra` file for editing.
 - b. Change the setting of the following property to true (it is false as shipped):
`TIBCO.BE.function.catalog.coherence=true`
 - c. Save the file and restart TIBCO BusinessEvents Studio.

For your information, log messages printed to the console at startup show the location of the files in use. See `BE_HOME/bin/logs/cep-engine.log`.

Configuring Oracle Coherence Cluster Discovery

See [Enabling Use of Oracle Coherence as the Cache Provider on page 433](#) for prerequisite actions you must take before you can use Oracle Coherence as the cache provider.

You can use either multicast or well-known-address (WKA) discovery, as appropriate. See Cache Cluster Member Discovery in *TIBCO BusinessEvents Architect's Guide* for basic guidelines. Procedures for both methods of discovery are explained below.

Guidelines for Managing Coherence Clusters

Provided Files The following files are located in `BE_HOME/lib/cep-datagrid-oracle.jar`:

- `coherence-cache-config-jdbc.xml` is an example cache configuration descriptor file.
- `tangosol-coherence-override-tibco-be.xml` is an example operational descriptor override file. To reference such a file, use the property `tangosol.coherence.override` in the CDD file.

To understand when and how to use them, read the Coherence documentation.

The following links to Coherence documentation provide helpful information for use of Coherence as the cache provider.

Checklist and Guidelines Before Architecting a New Project

http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/deploy_checklist.htm

http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/deploy_plat_consider.htm

http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/tune_perftune.htm

Coherence Network Protocol

http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/cluster_tcp.htm

http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/tune_datagramtest.htm

Coherence Metrics

http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/appendix_mbean.htm

http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/appendix_operational.htm

Configuring Multicast Cluster Discovery for Coherence Clusters

Multicast is the default option. If default values for multicast properties work in your environment, you can omit this procedure.

To Configure Multicast Cluster Discovery

1. If the default values for the following properties are not appropriate for your environment, add them to the properties sheet of the CDD Cluster tab and specify values as needed:

```
tangosol.coherence.clusteraddress
tangosol.coherence.clusterport
tangosol.coherence.ttl
```

See [Multicast Discovery Properties on page 438](#) for details

2. For multicast discovery you may also need to set these properties as explained in [Localhost and Localport Properties on page 439](#):

```
tangosol.coherence.localhost
tangosol.coherence.localport
```



Specifying one or more well-known addresses disables all multicast communication. Remove any well-known address properties, if any.

Configuring Well-Known Address Cluster Discovery for Coherence Clusters

You can add up to nine well-known addresses. If you need to add more refer to Coherence documentation for instructions about using override files.

To configure well-known address discovery, you must configure machine-specific settings at the cluster level and at the processing unit level.

One processing unit deployed to a WKA machine must have the additional WKA configuration. Additional processing units can be deployed to a WKA machine, configured in the usual way, and they will discover and join the cluster in the usual way at runtime.

To Configure Well-Known Address Cluster Member Discovery

For details about the properties, see [Well-Known Address Properties on page 441](#).

1. In the CDD file editor Cluster tab properties sheet, add a pair of WKA properties for each machine you want to configure as a well-known address machine:

```
tangosol.coherence.wkan HostIP
tangosol.coherence.wkan.port Hostport
```

For example at the cluster level you might have these two WKA machines:

```
Property: tangosol.coherence.wka1      Value: 10.97.118.151
Property: tangosol.coherence.wka1.port Value: 8098

Property: tangosol.coherence.wka2      Value: 10.97.118.152
Property: tangosol.coherence.wka2.port Value: 8098
```



If two engines run on one machine (with a single IP), ensure that each engine uses a different port. For example:

```
Property: tangosol.coherence.wka1      Value: 10.97.118.151
Property: tangosol.coherence.wka1.port Value: 8098

Property: tangosol.coherence.wka2      Value: 10.97.118.151
Property: tangosol.coherence.wka2.port Value: 8099
```

2. In the Processing Units tab properties sheet for a processing unit (PU), configure one set of WKA properties to match one of the cluster level set of WKA properties:

```
tangosol.coherence.localhost HostIP
tangosol.coherence.localport Hostport
```



If you will deploy using TIBCO BusinessEvents Monitoring and Management (MM) also add the following property:

```
be.engine.hostaddress HostIP
```

Use the same the value as the value of the `tangosol.coherence.localhost` property. See [Configuring for Coherence WKA Cluster Discovery](#) in Chapter 3, *Basic MM Configuration* in *TIBCO BusinessEvents Administration* for more steps required when configuring MM.

Set the `localhost` property to the IP of the host where you will deploy the PU, and set the `localport` property to the port defined in the cluster properties `localport` property.

3. Repeat [step 2](#) until you have configured a PU with matching PU-level properties for each cluster-level (that is, machine-level) set of WKA properties.

For example, if you configured the two well-known addresses shown in [step 2](#), then at the processing unit level you would configure two processing units as follows:

```
tangosol.coherence.localhost. Value: 10.97.118.151
tangosol.coherence.localport. Value: 8098

tangosol.coherence.localhost. Value: 10.97.118.152
tangosol.coherence.localport. Value: 8098
```

At deploy time you must deploy the processing units on the appropriate (matching) WKA machine. It can be helpful if the name of the PU contains the machine name as a reminder that this PU must be deployed to that machine.

See [Localhost and Localport Properties on page 439](#).

CDD Cluster Tab Coherence Properties Reference



Oracle Coherence Cache Provider This section is relevant only if you use Oracle Coherence as the cache provider.

TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

Add properties as needed to configure multicast cluster member discovery, or well-known address cluster member discovery. See [Configuring Oracle Coherence Cluster Discovery on page 434](#) for details.

Table 47 CDD Cluster Tab Coherence Properties (Sheet 1 of 6)

Property	Notes
Multicast Discovery Properties	
If you will define cluster members using multicast discovery properties, use the properties in this section, instead of those in the section Well-Known Address Properties on page 441 .	
<code>tangosol.coherence.clusteraddress</code>	<p>Use this setting if multicast discovery is used and if you need a non-default value. Specifies the multicast IP address that the socket will listen to or publish on.</p> <p>Possible values are addresses between (and including) 224.0.0.0 and 239.255.255.255.</p> <p>Default value is 224.3.3.1</p>
<code>tangosol.coherence.clusterport</code>	<p>Use this setting if multicast discovery is used and if you need a non-default value. Specifies the port that the socket will listen to or publish on.</p> <p>Possible values are integers between 1 and 65535.</p> <p>Default value is 35463</p>

Table 47 CDD Cluster Tab Coherence Properties (Sheet 2 of 6)

Property	Notes
<code>tangosol.coherence.ttl</code>	<p>Specifies the time-to-live setting for the multicast, that is, the maximum number of "hops" a packet can traverse. A hop is defined as a traversal from one network segment to another via a router.</p> <p>For production use, set this value to the lowest integer value that works. Setting the value too high can use unnecessary bandwidth on other LAN segments and can even cause the operating system or network devices to disable multicast traffic.</p> <p>On a single-host cluster, set to zero (0). On a simple switched backbone, set to 1. On an advanced backbone with intelligent switching, it may require a value of 2 or more.</p> <p>Note: A value of 0 is intended to keep packets from leaving the originating machine. However, some operating systems do not implement this correctly, and the packets may in fact be transmitted on the network.</p> <p>Required for multicast configuration.</p> <p>Possible values are integers between 0 and 255.</p> <p>Default value is 4</p>
<h3>Localhost and Localport Properties</h3>	
<p>These properties are used in these cases:</p> <ul style="list-style-type: none"> • When a host has multiple network cards. • For multicast discovery when more than one cluster is running on the same subnet (localhost is required but not localport in this case.) • Add these properties at the PU level when well-known address discovery is use as explained in Configuring Well-Known Address Cluster Discovery for Coherence Clusters on page 435. <p>Default values are provided at the cluster level. However if you need specify these properties at the PU level, add them as Processing Units tab properties and provide the values as needed.</p>	
<code>tangosol.coherence.localhost</code>	<p>Specifies the IP address that the socket will listen to or publish on.</p> <p>As needed, you can set the value of the <code>localhost</code> property to the value <code>localhost</code>. However, if <code>localhost</code> is used as the loop back address (127.0.0.1) you must enter a machine name or IP address.</p> <p>Default value is <code>localhost</code>.</p>

Table 47 CDD Cluster Tab Coherence Properties (Sheet 3 of 6)

Property	Notes
<code>tangosol.coherence.localport</code>	<p>Specifies the port that the socket will listen to or publish on.</p> <p>Possible values are 1 to 65535.</p> <p>Default value is 8088.</p> <p>Note If a specified port is not available, the object management layer (by default) increments the port number until it finds an available port. Avoid potential conflicts by choosing a number that is not close to a port used by other software in your environment.</p> <p>Tip To turn off auto-incrementing, add this property:</p> <pre>tangosol.coherence.localport.adjust=false</pre>

Table 47 CDD Cluster Tab Coherence Properties (Sheet 4 of 6)

Property	Notes
Well-Known Address Properties	
See Configuring Well-Known Address Cluster Discovery for Coherence Clusters on page 435.	
Note If you will discover cluster members using well-known addresses, use the properties in this section, and remove the multicast discovery properties shown in the section Multicast Discovery Properties on page 438.	
<code>tangosol.coherence.wkan</code>	
<code>tangosol.coherence.wkan.port</code>	
The addresses and ports for machines used by the well-known address cluster discovery protocol.	
At least one of these machines must be running at any time so that others can join the cluster.	
For <code>tangosol.coherence.wkan</code> , enter the IP address.	
For <code>tangosol.coherence.wkan.port</code> , enter a value between 1 and 65535.	
For example (in the UI the properties are not entered quite this way):	
<code>tangosol.coherence.wka1</code>	10.97.118.151
<code>tangosol.coherence.wka1.port</code>	8088
<code>tangosol.coherence.wka2</code>	10.97.118.152
<code>tangosol.coherence.wka2.port</code>	8088
Also at the Processing Units tab, configure <code>localhost</code> and <code>localport</code> properties for one processing unit that will be deployed to the WKA machine. Set the <code>localhost</code> value to the value of the <code>wkan</code> property. Set the <code>localport</code> value to the value of the <code>wkan.port</code> property. (See Localhost and Localport Properties on page 439)	
Note You can configure two well-known addresses for the same machine, and use a different port number for each. In this case you would also configure two processing units, each of which matches one set of WKA properties.	
Tip To turn off auto-incrementing, add this property:	
<code>tangosol.coherence.localport.adjust=false</code>	

Table 47 CDD Cluster Tab Coherence Properties (Sheet 5 of 6)

Property	Notes
Other Coherence Properties	
These properties are used in various situations.	
<code>tangosol.coherence.distributed.threads</code>	<p>Specifies the number of Coherence daemon threads used by the distributed cache service when Oracle Coherence is used as the cache provider.</p> <p>Mainly used with write-behind database strategy.</p> <p>When used with cache aside strategy, this setting is used for handling cache operations only (gets and puts).</p> <p>In this release the property must be set to the same value across the cluster.</p> <p>Value is per processing unit.</p> <p>A value of zero (0) means all relevant tasks are performed on the service thread.</p> <p>See Write Behind Options in <i>TIBCO BusinessEvents Architect's Guide</i>.</p> <p>Default value is 0</p>
<code>tangosol.coherence.override</code>	<p>Specifies the location of an Operational Descriptor Override File. A sample value is:</p> <pre>file:/c:/tmp/my_tangosol-coherence-override.xml</pre> <p>An example override file is provided in <code>BE_HOME/lib/cep-datagrid-oracle.jar</code>.</p> <p>Use of an operational descriptor override is not generally required. For details on override files, see Coherence documentation.</p>

Table 47 CDD Cluster Tab Coherence Properties (Sheet 6 of 6)

Property	Notes
<code>tangosol.coherence.localport.adjust</code>	<p>An auto-incrementing feature ensures that a different port is used if one specified is already in use. However in various situations you may want to turn off this behavior. For example, if you use TIBCO BusinessEvents Monitoring and Management, and the MM server runs on the same machine as any of the monitored cluster engines, you must explicitly ensure that all ports used by MM <i>and</i> the monitored cluster are unique. Therefore the auto-incrementing feature may not be appropriate.</p> <p>To turn off auto-incrementing, add this property and set the value to false.</p> <p>Default is true.</p>
<code>tangosol.coherence.guard.timeout</code>	<p>When infrastructure latency occurs such as remote databases, you can improve the performance of inference engines by setting the property to 0 in <code>be-engine.tra</code>.</p> <p><code>tangosol.coherence.guard.timeout=0</code></p>



TIBCO BusinessEvents Express This chapter relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

This chapter explains how to configure agents to work cooperatively as routers and receivers to ensure that related messages arriving from queue sources are handled by the same agent, so that related information is locally available.

Currently only queue messages from TIBCO Enterprise Message Service are supported for this configuration.

Topics

- [Understanding Load Balancing Options, page 446](#)
- [Content-Aware Load Balancer Configuration, page 449](#)
- [Adhoc Load Balancer Configuration, page 450](#)
- [CDD Load Balancer Tab Properties Reference, page 452](#)

Understanding Load Balancing Options

Load balancing is available for messages arriving from queues. Do not use load balancing for topic-based or other broadcast sources.

Two kinds of load balancing configuration are available: basic and content-aware load balancing. Both kinds support messages arriving from TIBCO Enterprise Message Service queue sources.

Every JMS destination that is configured to be an input destination runs in its own JMS Session. This provides good throughput on queues and topics for processing, and less connections.



The JMS message acknowledgement modes `EXPLICIT_CLIENT_ACKNOWLEDGE` and `EXPLICIT_CLIENT_DUPS_OK_ACKNOWLEDGE` are supported.

Basic Load Balancing

With basic load balancing, events from queue sources are automatically distributed between deployed instances of an agent class. To set up this kind of load balancing, you deploy multiple instances of an agent class that listens to a JMS destination. Each deployed agent instance runs in a different processing unit.

Certain aspects of the design have to be managed by the application. See [Designing for Concurrency on page 90](#) for related information.

This method can be useful when there is no relationship between the events that would require them to be processed in a certain order. If the order or grouping of events received is important, use content-aware load balancing. Content-aware load balancing has other benefits also, as explained below.

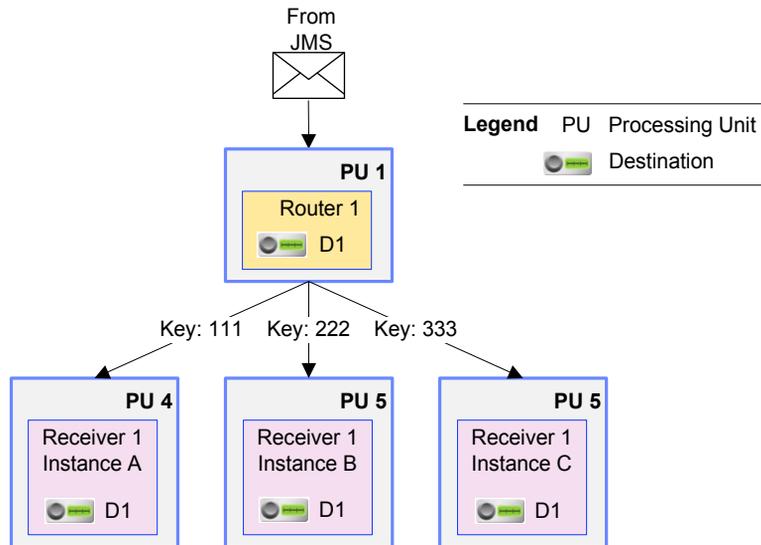
Content-aware Load Balancing

With content-aware load balancing, all related events arriving from queues are routed to the same agent using a routing key. The key uses the value of a selected event property. For example, if the event property is `ZipCode` then a routing key is a specific zip code. All messages relating to one zip code are routed (over TCP) to the same agent, providing "session stickiness."

Content-aware load balancing uses *routers* and *receivers*. You can configure routers and receivers from the CDD Load Balancer tab. One receiver can handle more than one set of related events. For example if the routing key is a zip code, one receiver might handle events for multiple zip codes.

Use of content-aware load balancing simplifies project configuration, and makes runtime behavior more efficient. For example, only local locking is generally required (whereas basic load balancing requires cluster-wide locking). Also the L1 cache does not have to be checked for version consistency.

The following diagram shows the CDD based router and receiver configuration.



Routers A router PU receives messages from the JMS server and routes them to appropriate receivers. Routers do no other work. For example, they should not execute rules.

A router PU contains one inference agent with one or more sets of JMS channels and destinations. Each destination has a default event. Values of one property of that event are used at runtime as *routing keys*.

Event preprocessors can be used as needed to populate the routing key property, for example using some calculation or combination of other event properties.

The router redirects events over TCP to a receiver, based on the destination and the routing key values. The router transparently distributes the load across the available receivers. If a receiver agent fails, its messages (that is, messages with the key that the router was sending to that agent) are routed to another receiver and continue to be handled by that other receiver.

Receivers Receivers are the inference or query agents that do the actual work. A receiver PU contains one inference or query agent. A set of receivers belongs to the same agent class. Receivers can also do other work, in addition to the work they receive from the router.

With CDD configuration, a receiver agent class is configured with one of the channel and destination configurations defined in the router. The destination, however, functions as a pseudo destination.

Content-Aware Load Balancer Configuration

When setting up content-aware load balancing, you first choose between two methods: pair configuration and adhoc configuration.

Pair Configuration

Preconfigured using the Load Balancer tab of the CDD.

The pair configuration uses a pair of processing units to act as routers and receivers. Two agent classes, a receiver inference class and a router inference class, are configured to use the same EMS destination. The router processes messages from the EMS queue and sends it to a receiver based on the routing key specified in the CDD.

See [CDD Load Balancer Tab Properties Reference on page 452](#) for details.

Adhoc Configuration

Allows minimal preconfiguration using the Load Balancer tab of the CDD. Use catalog functions to implement the load balancer at runtime.

Only the load balancer name and the local destination are configured in the CDD. The router agent and receiver agent classes are configured using catalog functions in their rules and rulefunctions.



Ensure that the destination has a default event associated with it.

Adhoc Load Balancer Configuration

This section provides guidelines to configure the router and receivers in an adhoc load balancer configuration.

Router Configuration

`LoadBalancer.Router.*` functions are used for the router side.

Create the Load Balancer

Use this rule function as a startup rule function. It creates and returns a load balancer that can be used to send messages to load balanced remote destinations. You create the router TCP connection in a startup rule function.

```
LoadBalancer.Router.createLoadBalancerTo(adhocConfigName);
```

Send Event to Receiver

Use this rule function as an event preprocessor. It sends an event to a remote receiver. The router agent does not have any destinations. The routing decision is made using the routing key.

```
void send(Object loadBalancer, SimpleEvent event, String routingKey)
```

Discard the Load Balancer

Put this rule function in a shutdown rule function.

```
Object loadBalancer = Util.HashMap.remove(String mapID, String key);
Util.HashMap.deleteMap(Object LoadBalancerReceiver );
```

This rulefunction discards the load balancer.

```
void discardLoadBalancer(Object loadBalancer)
```

Receiver Configuration

`LoadBalancer.Receiver.*` functions are used for the receiver side.

Local Channel

Receiver needs a local channel. The port information is obtained like this:

```
int port = System.getPropertyAsInt("receiver_localchnl_localdest_port",
34567);
```

Create a Receiver

Creates and returns a receiver object that receives messages from a router. Messages will be received from the router on the local channel and destination specified. Create the receiver TCP connection in a startup rule function.

```
Object LoadBalancer.Receiver.createTcpReceiverFor(String adhocConfigName)
```

Discard a Receiver

You discard the receiver in a shutdown rule function.

```
LoadBalancer.Receiver.discardReceiver(Object loadBalancedReceiver);
```

Receiver Membership Functions in Catalog

`LoadBalancer.Receiver.Membership.isInFlux`

Returns true if the loadbalancer node membership is in a state of change such as nodes joining and/or leaving currently or in the recent past.

`LoadBalancer.Receiver.Membership.getRecentChangeAt`

Returns the timestamp (1970 epoch milliseconds) at which the most recent membership change occurred.

CDD Load Balancer Tab Properties Reference

Configure the load balancer by specifying the receiver and router agents, and setting the destination.

Table 48 CDD Load Balancer Tab Properties

Property	Notes
Pair Configuration	
Name	Name of the pair configured load balancer.
JMS Destination	Destination used by the router and receiver agents. Ensure that the destination has a default event configured.
Key	Routing key used by the pair configuration.
Router	Router agent class for the load balancer configuration.
Receiver	Receiver agent class for the load balancer configuration.
Adhoc Configuration	
Name	Name of the adhoc load balancer. When creating a load balancer using the catalog function <code>createLoadBalancerTo</code> , the name of the load balancer must be specified. For example, <pre>Object loadBalancer = LoadBalancer.Router.createLoadBalancerTo("adhocLoadBalancerName");</pre>

Table 48 CDD Load Balancer Tab Properties

Property	Notes
JMS Destination	<p>Destination used by the router and receiver agents.</p> <p>In an adhoc configuration, a local channel with a local destination is used to communicate between the router and receiver agents.</p>
Properties	
Transport	<p>Transport to enable communication between the receiver and the router. Typically, this happens through an internal TCP connection.</p> <p>The default value is tcp.</p>
Hostname	<p>Hostname where the load balancer is configured.</p> <p>The default value is localhost.</p>
Port	<p>Port number used by the specified transport.</p>

Backing Store Configuration

This chapter explains how to configure CDD settings and properties that relate to backing store behavior.

Instructions for creating a backing store are in [Chapter 30, JDBC Backing Store Setup](#), on page 523.



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

Topics

- [Configuring Backing Store Settings and Properties](#), page 456
- [CDD Cluster Tab Backing Store Settings Reference](#), page 457
- [CDD Cluster Tab Backing Store Properties Reference](#), page 461
- [Setting up Shared Nothing Persistence](#), page 466

Configuring Backing Store Settings and Properties

This section provides a high level procedure and pointers to the reference tables and other sections that contain the details.

Oracle Database Strategy If the Strategy field for an Oracle database with a Shared All persistence (set under the Cluster tab > Backing Store) is set to `oracle`, then Oracle Database pooling strategy settings are used and various CDD properties act on the corresponding Oracle property, as noted in the reference table.



In some circumstances, it is necessary to configure some backing store settings before you set up the backing store. See [Cases That May Need Additional Setup on page 527](#) in [Chapter 30, JDBC Backing Store Setup](#).

To Configure the Cluster Tab Backing Store Settings and Properties

1. Open the CDD file you added, as explained in [Adding a Cluster Deployment Descriptor \(CDD\) on page 403](#).
2. Select the Cluster tab > Backing Store node on the left and on the right, configure settings as explained in [CDD Cluster Tab Backing Store Settings Reference on page 457](#).
3. Select the Cluster tab > Connection node on the left and on the right configure database connection settings. See [Backing Store > Connection Settings on page 459](#).
4. Select the Cluster tab > Properties node on the left and configure properties on the right, following guidelines in [CDD Cluster Tab Backing Store Properties Reference on page 461](#).
5. You can select domain objects (entities) to be included in or excluded from the backing store. In addition preloading options are available for loading domain objects from backing store to cache at system startup. See [Chapter 28, Domain Objects Configuration, on page 477](#).
 - See [Configuring Preloading Options on page 479](#) for object settings
 - See `be.engine.cluster.recovery.threads` in [CDD Cluster Tab Backing Store Properties Reference on page 461](#).
6. Save the resource.

CDD Cluster Tab Backing Store Settings Reference



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

For the related procedure, see [Configuring Backing Store Settings and Properties on page 456](#).

Table 49 CDD Cluster Tab Backing Store Settings (Sheet 1 of 4)

Property	Notes
Persistence Option: None	
	Specifies that the cluster does not have a backing store or the backing store is temporarily disabled, for example during testing phases of a project.
	Note: Individual entities can be set to not use the backing store. See Has Backing Store in CDD Cluster Tab Domain Object Override Settings Reference on page 485 .
	Default is set to None.
Persistence Option: Shared All	
Database Type	Select which of the supported DBMS products to use: Oracle, SQL Server, or Berkeley DB. Default is Oracle.
Strategy	Used only if the Database Type is set to Oracle. If you use Oracle Database, you have the option of using either the TIBCO BusinessEvents internal pooling implementation, or Oracle Database's implementation. Possible values are as follows: jdbc Use the internal pooling mechanism. oracle Use Oracle's pooling mechanism (see the class <code>OracleConnectionCacheManager</code> in the package <code>oracle.jdbc.pool</code>). When set to <code>oracle</code> then the TIBCO BusinessEvents pooling property values are used to set their corresponding to Oracle Database properties. Default is <code>jdbc</code> .

Table 49 CDD Cluster Tab Backing Store Settings (Sheet 2 of 4)

Property	Notes
Cache Aside	<p>Available only for the backing store database types Oracle and SQL Server.</p> <p>Choose between these two options:</p> <ul style="list-style-type: none"> <p>Checkbox unchecked means Write Behind Writes data to the cache and then to the backing store. One write-behind thread is used for each entity type.</p> <p>If write-behind strategy is used with Oracle Coherence cache provider, you can also set <code>tangosol.coherence.distributed.threads</code> in the cluster level properties.</p> <p>Checkbox checked means Cache Aside Writes data to the cache and at the same time to the backing store. User controls are available for threads and queue size, and other options such as using parallel or sequential operations in the post-RTC phase.</p> <p>See Chapter 10, Threading Models and Tuning in <i>TIBCO BusinessEvents Architect's Guide</i> for more information.</p>
Enforce Pools	<p>Available only for the backing store database types Oracle and SQL Server.</p> <p>Check this property if you want to enforce connection pool properties. See Database Connection Properties on page 461 for the property details.</p> <p>Default is unchecked.</p>
Berkeley DB Data Store Path	<p>Available only for the backing store database type Berkeley DB.</p> <p>Enter the absolute path of the data store to be used. The data store needs to be located on a network drive that can be accessed by all the cache nodes.</p>

Table 49 CDD Cluster Tab Backing Store Settings (Sheet 3 of 4)

Property	Notes
Persistence Option: Shared Nothing	
See Setting up Shared Nothing Persistence, page 466 for details on setting up Shared Nothing persistence.	
Persistence Path	<p>Specifies the absolute path to the directory where the data is to be stored. For example, <code>/tmp/datastore/</code>.</p> <p>Set this value for each individual processing units.</p>
Persistence Policy	<p>Specifies the type of communication to be used to maintain persistence: asynchronous (ASYNC) or synchronous (SYNC).</p> <ul style="list-style-type: none"> • ASYNC This policy is recommended if you want to avoid frequent IO operations, which can slow inference agents. • SYNC Solid State Drives (SSD) are recommended when using this policy. <p>Default is async.</p>
Backing Store > Connection Settings	
Try running with default pool values and monitor the behavior. Using more connections improves runtime performance and can also speed up recovery in the event of a failure.	
Pool settings are used only if Enforce Pools is checked.	
URI	<p>Specifies the project path, that is, the path from the project root to the JDBC Connection resource, to define the connection to the backing store. For example:</p> <pre>/SharedResources/JDBC Connection.sharedjdbc</pre> <p>You can also use a global variable to specify the connection.</p> <p>Default value is <code>%%DbUri1%%</code>.</p>

Table 49 CDD Cluster Tab Backing Store Settings (Sheet 4 of 4)

Property	Notes
Min Size	<p>Minimum number of JDBC connections in the JDBC connection pool used for the backing store.</p> <p>Oracle Database Strategy If the Cluster tab > Backing Store > Strategy field is set to <code>oracle</code>, then Oracle Database strategy settings are used, and this property corresponds to the <code>OracleConnectionCacheManager</code> class property <code>MinLimit</code>. Default is 10.</p>
Max Size	<p>Maximum number of JDBC connections in the JDBC connection pool used for the backing store. Connections do not exceed the maximum.</p> <p>The value of this property overrides the value of the Maximum Connections setting in the JDBC Connection resource.</p> <p>Although the limit is seldom reached, you can guarantee a connection is always available for a <code>dbwriter</code> thread as follows. Set the this field to the same value as the <code>Agent.AgentClassName.dbthreadcount</code> setting.</p> <p>Similarly (and also seldom needed), with Coherence cache provider, you can guarantee a connection is available by setting this field to the same value as the property <code>tangosol.coherence.distributed.threads</code>.</p> <p>Oracle Database Strategy If the Cluster tab > Backing Store > Strategy field is set to <code>oracle</code>, then Oracle Database strategy settings are used, and this property corresponds to the <code>OracleConnectionCacheManager</code> class property <code>MaxLimit</code>. Default is 10.</p>
Initial Size	<p>Specifies the initial size of the JDBC connection pool used for the backing store, when it is created on startup. For example:</p> <p>Oracle Database Strategy If the Cluster tab > Backing Store > Strategy field is set to <code>oracle</code>, then Oracle Database strategy settings are used, and this property corresponds to the <code>OracleConnectionCacheManager</code> class property <code>InitialLimit</code>. Default is 10.</p>

CDD Cluster Tab Backing Store Properties Reference



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

For the related procedure, see [Configuring Backing Store Settings and Properties on page 456](#).

Table 50 CDD Cluster Tab Backing Store Properties

Property	Notes
Database Connection Properties	
	Used only if Enforce Pools is checked (see CDD Cluster Tab Backing Store Settings Reference on page 457).
<code>be.backingstore.dburi.pool.waitTimeout.0</code>	<p>Used only if the Strategy setting (see CDD Cluster Tab Backing Store Settings Reference on page 457) is set to <code>oracle</code>.</p> <p>Oracle Database Strategy If the Cluster tab > Backing Store > Strategy field is set to <code>oracle</code>, then Oracle Database strategy settings are used, and this property corresponds to the <code>OracleConnectionCacheManager</code> class property <code>WaitTimeout</code>.</p> <p>Specifies behavior when a connection is requested and there are already Cluster tab > Backing Store > Connection > Max Size connections active. If the <code>be.backingstore.dburi.pool.waitTimeout.0</code> value is greater than zero (0), each connection request waits for up to the specified number of seconds. If no connection is returned to the pool before the timeout elapses, a <code>No Database Connection available</code> exception is thrown.</p> <p>The <code>waitTimeout</code> and <code>inactivityTimeout</code> properties specify wait periods to minimize the creation and destruction of connections (an expensive operation). Default is 1 second.</p>

Table 50 CDD Cluster Tab Backing Store Properties

Property	Notes
<code>be.backingstore.dburi.pool.inactivityTimeout.0</code>	<p>Oracle Database Strategy If the Cluster tab > Backing Store > Strategy field is set to <code>oracle</code>, then Oracle Database strategy settings are used, and this property corresponds to the <code>OracleConnectionCacheManager</code> class property <code>InactivityTimeout</code>.</p> <p>Specifies the number of seconds an unused connection remains available (so that other threads can use it). After this period, the connection is closed and removed from the pool.</p> <p>Default value is 900 seconds.</p>
<code>be.backingstore.readtimeout</code>	<p>Oracle Database Strategy If the Cluster tab > Backing Store > Strategy field is set to <code>oracle</code>, then Oracle Database strategy settings are used, and this property corresponds to the <code>OracleConnectionCacheManager</code> class property <code>ReadTimeout</code>.</p> <p>Use this property to handle situations where engines running inference agents hang when the JDBC connection to the backing store is slow or intermittent. If no response is received from the database within the specified period, a call is aborted.</p> <p>Time unit is milliseconds.</p> <p>A value of 0 (zero) means that no timeout is set.</p> <p>Default value is 0</p>
<code>be.backingstore.recreateOnRecovery</code>	<p>Set this property to true if the database pool size does not recover to the initial or minimum connection size, as defined by <code>Min Size</code> and <code>Max Size</code> properties (in CDD Cluster Tab Backing Store Settings Reference on page 457).</p> <p>Default value is false.</p>

Table 50 CDD Cluster Tab Backing Store Properties

Property	Notes
Other properties	
<code>be.backingstore.commitSize</code>	<p>Used with the Coherence cache provider and write-behind strategy only.</p> <p>Sets the maximum number of RTC transactions that a distributed cache service thread takes from the distributed cache service queue and processes in one batch. When threads are idle, they take jobs from the queue in smaller batches.</p> <p>Set this property to the desired number of transactions to suit your needs.</p> <p>See Write Behind Options in <i>TIBCO BusinessEvents Architect's Guide</i>.</p> <p>Default is 10.</p>
<code>be.backingstore.optimize.reads</code>	<p>Used with Microsoft SQL Server only.</p> <p>Set the property to true to improve the runtime performance.</p> <p>Use NOLOCK for SELECT statements to avoid locks on the database or table when SELECT statements are issued. An example syntax is:</p> <pre>select * from dbo.D_MailerIndex with (nolock) where ...</pre>
<code>be.backingstore.optimize.writes</code>	<p>Used with Microsoft SQL Server only.</p> <p>Set the property to true to improve the runtime performance.</p> <p>Use ROWLOCK with UPDATE or DELETE statements to avoid lock contentions. When you use rowlock in the T-SQL statement, the SQL Server locks only the rows that match the 'where' condition and not the entire table. An example syntax is:</p> <pre>DELETE FROM dbo.D_Mailed WITH (ROWLOCK) where mailernumber = '12345678895' and time_created\$ = 'somedate'</pre>
<code>be.backingstore.timestamp.useDataTimeZone</code>	<p>Used when the backing store is enabled.</p> <p>Set the property to true to ensure that the correct datetime properties are retrieved when an agent's timezone is changed and the agent is restarted.</p>

Table 50 CDD Cluster Tab Backing Store Properties

Property	Notes
<code>be.engine.cluster.cleanup</code>	<p>Used with write-behind strategy only. By default, deleted entities are removed from the backing store automatically at system startup. Set this property to <code>false</code> to disable that behavior. See Handling Entities Deleted from a Write-Behind Backing Store on page 555 for more details.</p> <p>Default is <code>true</code>.</p>
<code>be.engine.cluster.recovery.threads</code>	<p>Recovery threads are used when preloading the cache during startup.</p> <p>For an explanation of preloading and other preloading controls, see Configuring Domain Object Settings on page 478.</p> <p>Default is 4.</p>
<code>be.engine.cluster.recovery.distributed.batchsize</code>	<p>When distributed batch recovery is enabled (<code>be.engine.cluster.recovery.distributed.strategy=batch</code>), the recovery manager divides the target table into many smaller batches and assigns 1-to-n number of these batches to each node (for them to execute recovery).</p> <p>This parameter provides an approximation to the size of each such batch. Since batches are defined using the target table's key field (for example, approximated using ID\$ column), actual batch size will differ depending on the key distribution.</p> <p>Users can instead define <code>be.engine.cluster.recovery.distributed.batchpernode=2</code> and prevent that too few, or too many batches are created</p>

Table 50 CDD Cluster Tab Backing Store Properties

Property	Notes
<code>be.engine.cluster.useDBBatching</code>	<p>Note For use with cache aside and only when the parallel operations feature is used.</p> <p>This property has no effect if <code>Agent.AgentClassName.dbOpsBatchSize</code> is set to 1 (see Table 57, CDD Agent Classes Tab Inference Agent and Query Agent Properties, on page 501).</p> <p>This property affects how all RTC transactions that a database writer thread takes from the database operations queue are written to the backing store:</p> <ul style="list-style-type: none"> • When set to true, the RTC transactions are handled as one job. • When set to false, each RTC transaction is handled as a separate job. <p>For a guide to usage of this and other related properties, see the section Database Write Tuning Options for Cache Aside in Chapter 10, Threading Models and Tuning in <i>TIBCO BusinessEvents Architect's Guide</i>.</p> <p>Default is false.</p>

Setting up Shared Nothing Persistence

Shared Nothing persistence allows you to store data at individual node level, instead of a centralized location.



The following words are reserved words and cannot be the name of a property in an event or a concept:

`extid`, `id`, `closure`, `next`, `t1`, and `fired`.

Following steps explain how to configure a TIBCO BusinessEvents project to use Shared Nothing persistence:

1. In BusinessEvents Studio Explorer, edit the project CDD file to set the Object Management to **Cache** and Provider to **TIBCO**.



Shared nothing persistence works in the cache only mode. Any overrides for the concepts or events should also be in the cache only mode.

2. For the Backing Store, select the Persistence Option as **Shared Nothing**.
3. Enter the value for persistence path as the absolute path to the directory where data is to be stored.
4. Select the appropriate persistence policy.
5. Save the CDD file.

When using up Shared Nothing persistence, to ensure that data is not lost when nodes leave the cluster, you must set the number of backup copies to 1 or more.

Runtime Configuration to Specify the Engine Name Property

The TIBCO BusinessEvents engine name is used to name the Shared Nothing files for that node. To enable the engine name to be used, you need to explicitly set the engine name property when running clusters with Shared Nothing. The `-n <name>` must be specified for all the nodes (inference and cache) in the cluster.

Note that only one node will be started in the following cases:

- If the cache nodes in your cluster do not specify any names, the hostname of the machine on which the engine is running is used as the default engine name.
- If the cluster contains cloned cache nodes (nodes with the same engine name), only one cloned node will start.

Nodes with duplicate engine names will not be initialized.

Features Supporting the Shared Nothing Persistence Option

You can use the Shared Nothing persistence option with the following features:

- Concepts
- Events, with and without payloads
- Schedulers
- Scorecards
- State machines



When querying by extid, you must use the functions `Cluster.DataGrid.CacheLoadConceptByExtIdByUri` and `Cluster.DataGrid.CacheLoadEventByExtIdByUri`.

Chapter 27 **Berkeley DB Shared All Persistence**

This chapter explains how to configure the Berkeley DB Shared All Persistence option for persisting cache data to disk.

Topics

- [Berkeley DB Shared All Persistence Overview, page 470](#)
- [Configuring the Berkeley DB Shared All Persistence Option, page 472](#)
- [Reference to Berkeley DB Shared All Persistence CDD Properties, page 474](#)
- [Reference To Berkeley DB \(JE\) Properties, page 476](#)

Berkeley DB Shared All Persistence Overview

This form of data persistence uses TIBCO BusinessEvents DataGrid as the cache provider and Oracle Berkeley DB Java Edition software as the data store (not bundled with the TIBCO BusinessEvents software).



You must obtain a separate license for the Oracle Berkeley DB Java Edition software that is appropriate for your usage if you wish to use the Shared All with Berkeley DB feature.

With this form of persistence, cache servers act as persisters, and interact with the persistence layer. All cache nodes must have access to a reliable shared file system. The Berkeley DB data store is created within this shared file system.

The Berkeley DB Shared All Persistence option uses stores on a shared file system, such as NFS. This can result in better performance than use of a traditional DBMS product accessed over the network.

Reading from the Data Store

Persisters load data in bulk from the data store at startup. The following functions are used to load individual concepts into cache:

```
CacheLoadConceptByExtIdByURI()
```

```
CacheLoadConceptById()
```

Writing to the Data Store

Data is flushed to the operating system buffers for every write operation. TIBCO BusinessEvents uses `WRITE_NO_SYNC` durability by default, which means that Berkeley DB will flush every write to the operating system's buffers immediately but not call `fsync`. In case of an application crash, there will be no data loss as long as the underlying operating system synchronizes its buffers to disk.



Notes

- TIBCO BusinessEvents internally sets durability to `WRITE_NO_SYNC` durability. This behavior can be overridden by setting the `je.txn.durability` property in the `je.properties` file. See [Configuring the Berkeley DB Shared All Persistence Option on page 472](#).
- The `CacheLoadConceptByExtId()` function is not supported for use with this feature.
- Scorecards are not persisted (and so cannot be recovered).
- Scheduler events are not persisted (and so cannot be recovered)

Configuring the Berkeley DB Shared All Persistence Option

This section explains how to configure Berkeley DB Shared All Persistence. Configuration is done in the CDD file and optionally in a properties file called the `je.properties` file.

1. Download the supported version of the Oracle Berkeley DB Java Edition software from the Oracle web site. See the Release Notes for details on supported versions. Place the `je-versionNumber` JAR file in the following directory:

```
BE_HOME/hotfix/lib/ext/tpcl/
```

2. Open the project CDD or add one if the project does not yet have a CDD.
 - If you add a new CDD select Cache object management type in the wizard.
 - If you edit an existing CDD, select the Cluster tab. If the Object Management node is set to In Memory, right-click Object Management and select Cache.
3. Select the Cluster tab. In the Cluster tab navigation tree, select Object Management. In the Configuration panel, set the following:
 - Provider: TIBCO
 - Cache Agent Quorum: Generally set to the number of cache agents (see [Table 45, CDD Cluster Tab Cache OM Settings, on page 411](#) for details)
 - The following items are not relevant: Number of Backup Copies, Entity Cache Size, Object Table Cache Size.
4. In the Cluster tab navigation tree, select Backing Store. In the Configuration panel set the following values:
 - Persistence Option: **Shared All**
 - Database Type: **Berkeley DB**
 - Berkeley DB Data Store Path: The name and file location of the data store and (if used) the `je.properties` file.
5. If the CDD is using overrides for concepts or events, select the checkbox "Has Backing Store" to indicate that the entity will be stored in Berkeley DB. If the checkbox is not selected, then the entity will not be stored.
6. In the Cluster Tab navigation tree, select Properties. In the configuration tab, add the following properties (see [Table 51, Berkeley DB Shared All Persistence CDD Configuration Properties, on page 474](#) for details):

7. Select the Agent Classes tab. Configure cache agents with the properties shown in [Cache Agent Properties on page 474](#).
8. In the Agent Classes tab, configure inference agents with the properties shown in [Inference Agent Properties on page 474](#).
9. Save the CDD file and build project EAR files for deployment.
10. If you will use any properties shown in [Reference To Berkeley DB \(JE\) Properties on page 476](#), add the properties as name-value pairs in a file called `je.properties`. Place the file in the same location as the data store (as specified in the Berkeley DB Data Store Path setting).

Reference to Berkeley DB Shared All Persistence CDD Properties

This section provide a reference to the CDD properties

Table 51 Berkeley DB Shared All Persistence CDD Configuration Properties

Property	Notes
Cluster Settings	
Berkeley DB Data Store Path	<p>Specifies the name and file location of the data store. If this setting is not specified, a directory named <code>datastore</code> is created under the working directory.</p> <p>If you use a <code>je.properties</code>, file place it in the same location as the data store. See Reference To Berkeley DB (JE) Properties on page 476.</p>
Cache Agent Properties	
<code>be.backingstore.useobjecttable</code>	<p>Required.</p> <p>Required value is <code>false</code>.</p>
<code>be.engine.cluster.as.lock.ttl</code>	<p>Required. Value is specified in milliseconds. Controls the time after which TIBCO BusinessEvents DataGrid will forcibly unlock a space key. Thirty seconds is a reasonable value when the system not operating under heavy load (More than 2000 transactions per second may required a longer timeout period).</p> <p>Default value is 30000.</p>
Inference Agent Properties	
<code>be.backingstore.useobjecttable</code>	<p>Required.</p> <p>Required value is <code>false</code>.</p>
<code>be.engine.cluster.as.lock.ttl</code>	<p>Required. Value is specified in milliseconds. See notes for this property in the Cache Agent Properties section of the table.</p> <p>Default value is 30000.</p>

Table 51 Berkeley DB Shared All Persistence CDD Configuration Properties (Cont'd)

Property	Notes
<code>be.engine.cluster.as.node.retry.times</code>	<p>Required.</p> <p>Specifies the number of times TIBCO BusinessEvents retries a put or putAll call on the TIBCO BusinessEvents DataGrid cache.</p> <p>Each retry is done after five seconds. The number of retries depends on the <code>be.engine.cluster.as.lock.ttl</code> property. To calculate the value for retry times, use the following formula:</p> $\text{retry.times} = \text{lock.ttl} / 5 + 1$ <p>For example, if you set <code>be.engine.cluster.as.lock.ttl</code> to 30000, then you would set the <code>be.engine.cluster.as.node.retry.times</code> value to 7.</p>
<code>be.engine.cluster.as.log.dir</code>	<p>Directory where the log file is located, such as <code>C:/tibco/as/logs</code>.</p> <p>These ActiveSpaces log files are generated to allow that their parameters, such as their file name and size, are specified without filling up the system.</p>
<code>be.engine.cluster.as.log.filename</code>	<p>Name of the log file, such as <code>cs1.log</code></p>
<code>be.engine.cluster.as.logfile.count</code>	<p>Number of log files.</p>
<code>be.engine.cluster.as.logfile.size</code>	<p>Size of the log file</p>
<code>be.engine.cluster.as.logfile.append</code>	<p>Denotes whether the log file is to be appended.</p> <p>Options are true or false.</p>

Reference To Berkeley DB (JE) Properties

This section provides a reference to the Berkeley DB JE properties you can use to override the default configuration. For more details on the Berkeley DB JE properties, refer to the Berkeley DB product documentation:

<http://www.oracle.com/technetwork/database/berkeleydb/je-faq-096044.html>

http://docs.oracle.com/cd/E17277_02/html/GettingStartedGuide/administration.html#propertyfile

The properties are added to a file called `je.properties`. See [Configuring the Berkeley DB Shared All Persistence Option on page 472](#) for the required location.

Table 52 Berkeley DB JE Properties

Property	Notes
<code>je.txn.durability</code>	<p>Indicates the durability of a transaction. Possible values are <code>WRITE_NO_SYNC</code>, <code>NO_SYNC</code>, and <code>SYNC</code>.</p> <p>TIBCO BusinessEvents internally sets durability to <code>WRITE_NO_SYNC</code>. You can override that setting using this JE property.</p>
<code>je.log.numBuffers</code>	<p>The number of write buffers that the Berkeley DB uses internally.</p> <p>Default value is 3.</p>
<code>je.log.bufferSize</code>	<p>The size of each write buffer.</p> <p>Default value is 1048576.</p>
<code>je.log.totalBufferBytes</code>	<p>The sum of sizes of all the write buffers.</p> <p>Default value is 3145728.</p>
<code>je.maxMemoryPercent</code>	<p>The percentage of the JVM maximum memory to which the cache is limited.</p> <p>Default value is 60%</p>

This chapter explains options that determine how entity object instances are stored and options for preloading them from backing store to cache at system startup.



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

Topics

- [Configuring Domain Object Settings, page 478](#)
- [CDD Cluster Tab Domain Objects Default Settings Reference, page 481](#)
- [CDD Cluster Tab Domain Object Override Settings Reference, page 485](#)

Configuring Domain Object Settings



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.



If a backing store is already set up, and you enable any Use Backing Store settings, you must update the backing store setup. The backing store will not operate correctly unless you do so. See [Updating an Existing Backing Store Schema on page 552](#).

Domain object settings let you configure various behaviors for objects generated by the inference engines and stored in a cache. Many options relate to the way objects move between cache and backing store, so that you can tune memory usage and performance as needed.

You can configure the various behaviors globally (at the default level) and you can set overrides at the object type level. (Not all object level settings, however, are overrides.)

The main options are as follows. (Other options pertain to more specific situations and all are documented in the reference tables.)

- The mode: Cache plus Memory, Cache Only, or Memory Only. See Chapter 8, Cache Modes and Project Design in *TIBCO BusinessEvents Architect's Guide* to understand the effect of the different modes. Cache plus Memory (Cache + Memory in the UI) is a deprecated feature.
- Whether the objects or handles or both are preloaded into the backing store at startup.
- A preload fetch size (one setting for both objects and handles).
- Whether the cache is limited or unlimited. If limited, you can specify the size of the cache. See [Configuring a Limited \(or Unlimited\) Cache on page 408](#).

At the individual object type level only, you can also configure the following:

- Whether the object is stored in the backing store or not.
- A subscription preprocessor (used for Cache Plus Memory mode only).
- A backing store table name (used when setting up a backing store. See [Chapter 30, JDBC Backing Store Setup, on page 523](#)).

The settings are applied at the object level. For example, a contained concept can have a different limited cache setting from its container concept, and could be evicted from the cache at a different time.

Configuring Preloading Options

Preloading refers to the loading of the cache with objects from the backing store, at system startup. If you do not preload, then the objects are loaded as needed, which can affect performance the first time each object is requested, depending on the size of the objects.

At the default level, you can choose to preload or not preload two types of items separately: the objects themselves, and the handles to the objects, which are stored in a separate table called the object table. You can also specify the number of objects to preload (one setting for both types of items).

At the individual object level, you can override the preload setting as follows:

- Use the default setting
- Preload (True)
- Do not preload (False)

You can also specify (or override) how many objects to preload.



One tuning property for preloading is available:

```
be.engine.cluster.recovery.threads
```

To Configure the Cluster Tab Domain Object Settings and Properties

1. Add a CDD as explained in [Adding a Cluster Deployment Descriptor \(CDD\) on page 403](#), or open an existing CDD, configured for Cache OM.
2. Select the Cluster tab > Domain Objects > Default node on the left and on the right, configure settings as explained in [CDD Cluster Tab Domain Objects Default Settings Reference on page 481](#).
3. To add object-level overrides and other object-level settings do the following:
 - a. Select the Cluster tab > Domain Objects > Overrides node on the left and click **Add**.
 - b. At the Entity Selection dialog, select the ontology object type or types you want to configure and click **OK**. The first entity in the list is selected and the configuration section for the entity displays on the right.
 - c. Select the `/uri` node for each selected entity in turn and configure the settings on the right as needed. You can also edit existing override entries,

and remove entries not needed (by clicking Remove). See [CDD Cluster Tab Domain Object Override Settings Reference](#) on page 485.



When a project is migrated from 4.x to 5.1, domain object override entries are automatically added for all entities. See Migration from 4.x to 5.1 in *TIBCO BusinessEvents Installation* for details.

4. Save the resource.

CDD Cluster Tab Domain Objects Default Settings Reference

The Domain Object Default settings apply to all objects except those for which you explicitly configure overrides, using the Domain Object Overrides section (see [CDD Cluster Tab Domain Object Override Settings Reference on page 485](#)).

Table 53 CDD Cluster Tab Domain Object Default Settings (Sheet 1 of 4)

Property	Notes
Mode	<p>With Cache OM, you can keep memory objects in the cache or Rete network using the following cache modes.</p> <ul style="list-style-type: none"> • Memory Only: Objects are not persisted in the cache. They are kept in the Rete network (working memory) only. • Cache Only: Objects are persisted in the cache. They must be loaded into working memory as needed. This is the most common choice for a cache cluster. • Cache+Memory (also written as Cache Plus Memory): Deprecated feature. Due to issues with concurrency, in clusters with multiple active inference agents, use Cache Plus Memory only for constants and objects that change infrequently. <p>Note If you set the mode to Memory Only, the rest of the properties in this section are not relevant and are ignored.</p> <p>See Chapter 8, Cache Modes and Project Design in <i>TIBCO BusinessEvents Architect's Guide</i> to understand the effect of this setting.</p> <p>Default is Cache Only.</p>
Preload Entities	<p>Specifies whether objects are loaded into the cache from the backing store on system startup (both normal startup and recovery).</p> <p>Checked: All objects are preloaded into the cache from the backing store. Lower level settings can override this setting by excluding specified objects.</p> <p>Unchecked: No objects are preloaded. into the cache from the backing store. Lower level settings can override this setting by including specified objects.</p> <p>Default is unchecked.</p>

Table 53 CDD Cluster Tab Domain Object Default Settings (Sheet 2 of 4)

Property	Notes
Preload Handles	<p>Specifies whether object handles are loaded into the <code>ObjectTable</code> cache. The <code>ObjectTable</code> cache holds references (handles) to the objects themselves.</p> <p>Handles are used in the object table. See <i>The Role of the Object Table in TIBCO BusinessEvents Architect's Guide</i> for more details.</p> <p>Checked: All object handles are preloaded. Lower level settings can override this setting by excluding handles for specified objects.</p> <p>Unchecked: No object handles are preloaded into the cache from the backing store. Lower level settings can override this setting by including handles for specified objects.</p> <p>Default is unchecked.</p>
Preload Fetch Size	<p>If <code>Preload Entities</code> or <code>Preload Handles</code> or both are checked, this setting specifies the number of entity objects or handles (or both) to preload for each entity type whose objects or handles (or both) are configured to be preloaded.</p> <p>This setting applies to both objects and handles and cannot be set differently for each.</p> <p>Objects and handles are fetched in a non-deterministic manner.</p> <p>This setting can be overridden at the entity level.</p> <p>Set to 0 to preload all. Set to a number to load that number of objects or handles (or both).</p> <p>Default is 0. Ignored unless <code>Preload Entities</code> or <code>Preload Handles</code> or both are checked.</p>

Table 53 CDD Cluster Tab Domain Object Default Settings (Sheet 3 of 4)

Property	Notes
Check for Version	<p>This field applies to concepts that use cache-only mode or cache+memory mode.</p> <p>An inference agent uses its L1 cache, a local cache of limited size, to improve access time to the concepts stored in the cluster cache. When an agent finds a concept instance in this local cache, the Check for Version setting determines whether the agent will use the instance directly, or instead check in the cluster cache for more recent version.</p> <p>If Checked (default value) The agent checks in the cluster cache for a more recent version. If a more recent version exists, it will be used, and will replace the one found in the local cache.</p> <p>If Not Checked The agent uses the instance found locally.</p> <p>When content-aware load balancing is used, the local instance can be used without checking for version, improving performance.</p> <p>Default is checked.</p>
Constant	<p>This field applies to entities that use the cache-only mode or the cache+memory mode. In multi-engine deployments, use only for entities that do not change after creation.</p> <p>The processing unit has a special local cache used only for entities marked as Constant. Entities placed in this cache are only removed when they are explicitly deleted. If the processing unit finds an entity in the constant cache, it will use it without checking in the cluster.</p> <p>If Checked The entity is marked <code>Constant</code>, and uses the constant cache.</p> <p>If Not Checked The entity does not use the constant cache.</p> <p>Default is unchecked.</p>

Table 53 CDD Cluster Tab Domain Object Default Settings (Sheet 4 of 4)

Property	Notes
Evict from Cache on Update	<p>Used only if both of the following are the case:</p> <ul style="list-style-type: none"> • Cache-aside database write strategy is used • The property <code>Agent.AgentClassName.cacheTxn.updateCache</code> is set to false (see CDD Agent Classes Tab Properties Reference on page 501) <p>If checked: When a rule action changes the value of any of an entity's properties, then the entity instance is evicted from the cache (updates are saved in the backing store)</p> <p>Use as needed to improve performance and cache memory management. For example, if an entity is not accessed frequently, it may save memory in the cache if the entity is evicted from cache after it is updated.</p> <p>Possible values are checked (true) and unchecked (false).</p> <p>Default is unchecked.</p>
Is Cache Limited	<p>If checked, the cache size is limited.</p> <p>Limited cache requires use of a backing store. See Configuring a Limited (or Unlimited) Cache on page 408.</p> <p>The size of the entity cache and the size of the object table cache are set in the Object Management section of the Cluster tab.</p> <p>If not checked, the cache size is unlimited.</p> <p>You can override this default setting in entity overrides.</p> <p>Default is unchecked.</p>
Subscribe Cluster	<p>For objects that use Cache+Memory mode, check this checkbox to subscribe to subscription RTCs, so that changes to this object in one Rete Network are also changed in the Rete networks across all inference agents. See Using Locks to Ensure Data Integrity Within and Across Agents in TIBCO BusinessEvents Architect's Guide for details.</p> <p>Default is checked.</p>

CDD Cluster Tab Domain Object Override Settings Reference

Many settings simply override the value of default settings. See [Table 53, CDD Cluster Tab Domain Object Default Settings, on page 481](#) for general details about the use of each setting. The table below only provides details that are specific to overrides.

Table 54 CDD Cluster Tab Domain Object Override Settings (Sheet 1 of 4)

Property	Notes
Entity URI	<p>Specifies the project path to the entity for which overrides are being set. Defaults to the selected entity's URI. For example: <code>/Concepts/MyConcept</code>.</p>
Mode	<p>Overrides the Default level setting for this object type.</p> <p>Memory Only Mode If you set the mode for an entity to Memory Only, the rest of the properties in this section are not relevant and are ignored. Backing store is disabled for entities that use Memory Only mode.</p> <p>Caution If you change from Memory Only mode to a cache mode after the backing store has been set up, you must update the backing store schema. See Updating an Existing Backing Store Schema on page 552.</p>
Preload Entities	<p>Specifies whether objects of the specified type are loaded into the cache from the backing store on system startup (both normal startup and recovery).</p> <p>Overrides the Preload Entities setting at the Default level.</p> <p>default: Use the Preload Entities setting specified at the default level.</p> <p>true: Objects of the specified type are preloaded into the cache from the backing store. If the default level setting is not to preload entities, you can use this override to preload selected entities.</p> <p>false: No objects of the specified type are preloaded into the cache from the backing store. If the default level setting is to preload entities, you can use this override to not preload selected entities.</p> <p>Default is default.</p>

Table 54 CDD Cluster Tab Domain Object Override Settings (Sheet 2 of 4)

Property	Notes
Preload Handles	<p>Specifies whether object handles for the specified type are loaded into the cache from the backing store on system startup (both normal startup and recovery).</p> <p>Overrides the Preload Handles setting at the Default level.</p> <p>default: Use the Preload Handles setting specified at the default level.</p> <p>true: Handles for the specified type are preloaded into the cache from the backing store. If the default level setting is not to preload handles, you can use this override to preload selected entities' handles.</p> <p>false: No handles for the specified type are preloaded into the cache from the backing store. If the default level setting is to preload handles, you can use this override to prevent preloading the selected entities' handles.</p> <p>Default is default.</p>
Preload Fetch Size	<p>Overrides the Preload Fetch Size setting in the Default settings.</p>
Check for Version	<p>Overrides the value of the same-named setting in the Default settings.</p>
Constant	<p>Overrides the value of the same-named setting in the Default settings.</p>
Evict from Cache on Update	<p>Overrides the value of the same-named setting in the Default settings.</p>
Is Cache Limited	<p>Overrides the value of the same-named setting in the Default settings.</p>
Subscribe Cluster	<p>If this object uses Cache+Memory mode, check this checkbox to subscribe to subscription RTCs, so that changes to this object in one Rete Network are also changed in the Rete networks across all inference agents. See <i>Using Locks to Ensure Data Integrity Within and Across Agents</i> in <i>TIBCO BusinessEvents Architect's Guide</i> for details.</p>

Table 54 CDD Cluster Tab Domain Object Override Settings (Sheet 3 of 4)

Property	Notes
Subscription Preprocessor	<p>If this object uses Cache+Memory mode, and Subscribe Cluster is checked, specify a subscription preprocessor. This preprocessor is generally used to provide locking to ensure data consistency.</p> <p>The required signature for a subscription preprocessor is as follows:</p> <pre>boolean FunctionName(long id, String extId, int typeId, int version, boolean isDeleted)</pre> <p>See <i>Using Locks to Ensure Data Integrity Within and Across Agents in TIBCO BusinessEvents Architect's Guide</i> for details.</p> <p>Cache plus memory is a deprecated feature. See <i>TIBCO BusinessEvents Release Notes</i> for details.</p>
Backing Store Section	
Has Backing Store	<p>Used only if the Backing Store > Enabled checkbox is checked. To exclude an entity from the backing store, uncheck the Has Backing Store checkbox.</p> <p>Caution If you enable this override setting after the backing store has been set up, you must update the backing store schema. See Updating an Existing Backing Store Schema on page 552.</p> <p>Concepts Related by Containment or Inheritance All concepts related by containment or inheritance must either be included in the backing store or excluded from the backing store. That is, they must share the same value for the Has Backing Store setting.</p> <p>Default value is checked.</p>
Table Name	<p>Specifies a table name to be used in the backing store. Typically used if the entity name is long. See Ontology Identifiers that Exceed the DBMS Maximum Column Length on page 527 for details.</p>
Properties Metadata Section	
Property	<p>Displays the property name. Read-only.</p>

Table 54 CDD Cluster Tab Domain Object Override Settings (Sheet 4 of 4)

Property	Notes
Present in Index	<p>If checked, sets an unordered index on the property.</p> <p>Used only if Oracle Coherence is the cache provider. Also used only if the query agent has enabled indexing: see <code>be.agent.query.enable.filter.optimizer</code> in CDD Agent Classes Tab Properties Reference on page 501. Also see Indexing for More Efficient Cache Queries on page 311 for more details and options.</p>
Max Length	<p>Used with backing store to specify the length of string properties that exceed 255 characters (that is the actual contents stored in the column is more than 255 characters). Specifies the expected maximum length for the property. See String Properties that Exceed the DBMS Maximum Column Length on page 528 for details.</p>
Reverse References	<p>This setting is for use only with <code>ConceptReference</code> type concept properties.</p> <p>With a backing store, database updates related to a referring concept in a referenced concept can cause decreased performance. This happens when there are very many reverse references in a shared instance (referenced by many other instances).</p> <p>To address this issue, set the value to <code>false</code> for <code>ConceptReference</code> type properties.</p> <p>If you set the value to <code>false</code>, you must explicitly remove <code>ConceptReference</code> properties for deleted referenced concepts in the referring concept in your code.</p> <p>For example, if <code>employee</code> is a <code>ConceptReference</code> type property in a concept <code>acme</code>, and <code>smith</code> is an instance of a concept type <code>employee</code>, then you would set <code>ReverseReferences</code> to <code>true</code> for the <code>employee</code> <code>ConceptReference</code> property, and you would add something like this to rules:</p> <pre>acme.employee = null; Instance.deleteInstance(smith);</pre> <p>Or, for array properties:</p> <pre>Instance.PropertyArray.removeConceptReference(acme.employee, smith); Instance.deleteInstance(smith);</pre> <p>Default is <code>true</code>.</p>

Chapter 29 **Agent and Processing Unit Configuration**

This chapter explains how to configure agent classes and processing units in the CDD file.

Topics

- [Collections, Agent Classes, and Processing Units, page 490](#)
- [Configuring Collections of Rules, Rule Functions, and Destinations, page 492](#)
- [CDD Collections Tab Input Destination Settings Reference, page 494](#)
- [Configuring Agent Classes \(All OM Types\), page 496](#)
- [CDD Agent Classes Tab Settings Reference, page 498](#)
- [CDD Agent Classes Tab Properties Reference, page 501](#)
- [Understanding and Configuring Log Configurations, page 506](#)
- [CDD Collections Tab Log Configurations Settings Reference, page 509](#)
- [Logging for TIBCO BusinessEvents DataGrid on page 512](#)
- [Configuring Processing Units \(All OM Types\), page 513](#)
- [CDD Processing Units Tab Settings Reference, page 515](#)
- [CDD Processing Units Tab Coherence Log Properties Reference, page 517](#)
- [CDD Processing Units Tab JMS Server Connection Properties, page 520](#)

Collections, Agent Classes, and Processing Units

Destinations require additional configuration, which can be done in this tab. (Destinations that are added to agent classes individually can be configured at the Agent Classes tab.)

Collections

At the Collections tab, you can (optionally) group rules, rule functions, and destinations into collections so that they can be easily assigned to agent classes (and processing units in the case of log configurations).

See [Configuring Collections of Rules, Rule Functions, and Destinations on page 492](#).

Agent Classes

Agent classes define the different sorts of agents you can deploy. Various agent types are available depending on the object management (OM) type and on the add-on products used. Each agent type is configured differently.



See [Using Properties at Different Levels on page 400](#) to understand the effect of using agent class properties at the cluster level and at the processing unit level to widen the scope of the property.

See [Configuring Agent Classes \(All OM Types\) on page 496](#)

How Collections and Individual Resources are Used to Configure Agents

Different agent types use different types of resources.

In the Agent Classes node (on the left side of the CDD editor) you see categories of collections. Here, you add collections you defined earlier, as needed to configure the agent class:

Rules (Inference agent classes only.) It can be convenient to organize rules into collections for use in different inference agent classes. Select rule collections and individual rules as needed to define what rules will execute on inference agents of such classes at runtime.

Input Destinations Different agents listen for messages arriving at different destinations. When you select a destination for use in a collection or an individual agent, you add deploytime configuration settings, to create a *destination configuration*. For example, you define an event preprocessor and a threading model to use. Each destination configuration is assigned a unique ID.

Startup Functions and Shutdown Functions Select function collections and individual functions as needed, to define which functions execute at engine startup and shutdown respectively.



How Startup Rulefunctions and Shutdown Rule Functions are Executed

- The order of the functions (including the order of functions within collections) is the order in which they execute at runtime.
- Put startup rule functions (for use at start up) into different collections from shutdown rule functions (those used at shut down) so you can select them appropriately at the agent classes tab.

Log Configurations

Also in the Collections tab, you can add different log configurations. These are used when you configure processing units.

See [Understanding and Configuring Log Configurations on page 506](#)

Configuring Collections of Rules, Rule Functions, and Destinations

The purpose of configuring collections of rules, rule functions, and destinations, is to make it simpler to configure agent classes. When you configure an agent, you can add collections of resources or individual resources or both. Two collections are predefined: an all-rules collection and an all-functions collection.

Using References

A collection can have references to items (rules, rule functions, or destinations), and also references to other collections of the same type. References are identified in the groups tree by a reference symbol (↗). This mechanism enables you to reuse collections for more efficient configuration.

To Configure Collections

The procedure is in general the same for rules, rule functions, and destinations, so in these instructions, the word *item* is used to refer to any rule or rule function or destination.

For the log configurations procedure, see [To Add a Log Configuration on page 507](#)

1. In the Collections tab do any of the following:
 - To add a new collection, select the parent for the collection type, Rules, Destinations, Functions, or Log Configurations as needed, and click **Add**.
 - In the *Item* Collection field that appears on the right, enter a name for the group and click **Add** again.

Ensure that the collection name is unique across different collections in the CDD. For example, Rule collections and Destination collections in the CDD cannot have the same name.
 - To add *items* and *item* group references to a collection select the item collection and then click **Add**.

You see the Select *Items* dialog.
2. In the Select *Items* dialog do any of the following:
 - To add *items*, in the **Items** tree click the checkboxes of *items* you want to add to the group you are defining.
 - To add collection references, in the **Collection References** tree click the checkboxes of collections you want to add (by reference) to the collection you are defining.

When you select a collection on the left, you see details on the right: For example, the path to item you selected, and the names of collections you selected.

3. For function collections only, reorder the functions as needed, so that they execute in the correct order at runtime (that is, at startup or shutdown). Highlight a rule function in the tree on the left, and then click Move Up or Move Down as needed.
4. For destination collections only, configure each destination in turn. Select the destination on the left and complete the settings on the right to define characteristics such as the threading model to use, and the event preprocessor. See [CDD Collections Tab Input Destination Settings Reference on page 494](#) for information about each setting.
5. Save the CDD.

To Update Collections

- To remove an item in a collection or the collection itself, select the item or the group on the left and click **Remove**.
- To reorder rule functions in a function collection, select a rule function in the tree on the left, then click Move Up or Move down. This is important for startup and shutdown rule functions. Ensure that you put startup and shutdown rule functions into appropriate separate collections.
- You can change the URI (project path) of project resources to match their actual locations. To change the URI of an item, select the item on the left and edit the URI on the right.

CDD Collections Tab Input Destination Settings Reference

These agent-specific destination configuration settings are also available from the Agent Classes Tab. Collections enable you to configure once and use in multiple agents..

Table 55 CDD Collections Tab Input Destination Settings (Sheet 1 of 2)

Property	Notes
Input Destination ID	<p>Uniquely identifies this destination configuration at runtime. Edit as needed to ensure that each destination in the cluster has a unique deployment name.</p> <p>Default value is destination name.</p>
URI	<p>Project path to the destination (that is path to the destination in the design-time project).</p>
Preprocessor	<p>Specifying an event preprocessor for a destination is optional.</p> <p>Tip If you specify a preprocessor, you generally also also specify worker thread settings, because event preprocessors are multithreaded (unless Caller's Thread threading model is used, which is single threaded).</p> <p>Select the rule function that has been configured as this destination's event preprocessor.</p> <p>Event preprocessors are rule functions with one argument of type simple event. See Event Preprocessors on page 275.</p>

Table 55 CDD Collections Tab Input Destination Settings (Sheet 2 of 2)

Property	Notes
Threading Model	<p>If you specified a preprocessor, also specify thread settings. Select one model:</p> <p>Shared Queue Uses the TIBCO BusinessEvents system-wide shared queue and threads. For queue size and number of threads settings, see CDD Agent Classes Tab Settings Reference on page 498.</p> <p>Caller's Thread Uses the thread (and queue size) provided by the channel resource client. There is one thread per destination.</p> <p>Note If it is important to ensure that acknowledgements are sent in the expected order with Caller's Thread threading model, do not use the parallel operations option. See <code>Agent.agentClassName.enableParallelOps</code> in Table 57, CDD Agent Classes Tab Inference Agent and Query Agent Properties, on page 501.</p> <p>Destination Queue TIBCO BusinessEvents creates a dedicated thread pool and set of worker threads in each destination. See Thread Count and Queue Size below.</p> <p>For more information on threading models see Chapter 10, Threading Models and Tuning, in <i>TIBCO BusinessEvents Architect's Guide</i>.</p>
Thread Count	<p>If you specified Destination Queue in the Threading Model setting, specify the number of threads for this destination here.</p>
Queue Size	<p>If you specified Destination Queue in the Threading Model setting, specify the queue size for this destination here.</p>

Configuring Agent Classes (All OM Types)



TIBCO BusinessEvents Express Content relating to Cache OM and backing store does not apply to TIBCO BusinessEvents Express edition.

Agent class types are as follows:

Inference Agent Used with all OM types. For inference agent classes, you distribute a project's resources among the agent classes to define the specific work each agent will do. In Memory OM uses only inference agents, and each agent operates independently. With Cache OM, the agents share the cache data (as explained in *TIBCO BusinessEvents Architect's Guide*).

Cache Agent Used with Cache OM only. Cache agents handle distributed cache object storage. A processing unit can contain only one cache agent.

Query Agent Used with Cache OM only and available only if TIBCO BusinessEvents Event Stream Processing software is used.

Dashboard Agent Used with Cache OM only and available only if TIBCO BusinessEvents Views software is used. See *TIBCO BusinessEvents Views Developer's Guide* for details.

Monitoring & Management Agent Used internally by the Monitoring and Management component. Do not add any agents of this class. Configuration of this agent type is explained in Chapter 3, Basic MM Configuration of *TIBCO BusinessEvents Administration*.

Process Agent Used with Cache OM only and available only if TIBCO BusinessEvents Process Orchestration software is used.

To Add an Agent Class

You can begin by configuring classes provided by the wizard. You can rename the classes as desired. Then add more classes as needed.

1. In the Agent Classes tab, click **Add Agent**.
2. In the New Agent Class dialog enter an Agent Class Name and select an Agent Class Type from the list (see section introduction). Valid types for your project depend on object management type, and whether you use any TIBCO BusinessEvents add-on products. Click **OK**.

3. The new agent class name appears on the left. Select the agent class name. Appropriate settings for that agent type appear in the Configuration section.
 - Complete the settings as explained in [CDD Agent Classes Tab Settings Reference on page 498](#)
 - Add any properties as needed. The available properties are explained in [CDD Agent Classes Tab Properties Reference on page 501](#).
4. For inference, dashboard, and query agent class types, add the resources you want to use. In the agent class tree on the left, click each type of resource collection in turn and configure as explained next.

(In the instructions below, the word *item* stands in for destination, function, and rule depending on the collection category.)

 - a. Highlight a category of collections (for example Input Destination Collections).
 - b. Click **Add**. You see the Select *items* dialog.
 - c. In the upper section of the dialog, select individual project *item* resources, as desired.
 - d. In the lower section of the dialog (the Reference Groups section), select *item* collections you defined earlier, as desired.
 - e. Click **OK**. A list of *item* IDs appears in the box on the right.
5. If you added any individual destinations to the Input Destinations Collections category, highlight their name on the left and configure their settings on the right. See [CDD Collections Tab Input Destination Settings Reference on page 494](#) for details.

(Destinations within input destination collections are configured at the Collections tab.)
6. Do any of the following as needed:
 - Click a collection category on the left to see a list of collections and *items* you selected from that category on the right.
 - Expand a category on the left and click a collection reference within it. You see a list of its item IDs and paths, and any collection references within that collection, on the right.
 - Edit the project paths for individual items you add here. You would do this only if the project location of that item changed.
7. Save the CDD.

CDD Agent Classes Tab Settings Reference



TIBCO BusinessEvents Express Content relating to Cache OM and backing store does not apply to TIBCO BusinessEvents Express edition.

The tables below explain settings used with inference agents and query agents. Query agents are used with the TIBCO BusinessEvents Event Stream Processing add-on.

Dashboard agents are used with the TIBCO BusinessEvents Views add-on and are documented in *TIBCO BusinessEvents Views Developer's Guide*.

Table 56 CDD Agent Classes Tab Inference Agent and Query Agent Settings (Sheet 1 of 3)

Setting	Notes
Inference Agent and Query Agent Settings	
Max Size (Local Cache)	<p>Specifies the maximum number of objects (entities) in each agent's L1Cache (inference agent) or local cache (query agent). The L1 cache is a local cache used by the inference agent for local access to recently used objects. It is used to optimize access to objects.</p> <p>The query local cache is used in a way similar to the inference agent L1Cache. The query agent's local cache stores cache data locally for efficient reuse. The local cache listens to and synchronizes the locally stored entity instances with those in the main cache, so that the local cache stays up-to-date.</p> <p>When the threshold is reached, oldest entities are removed first.</p> <p>Default is 1024 (unit is objects).</p>
Eviction Time (Local Cache)	<p>Specifies an age limit on the cached entities in seconds. After this period, they are removed from the local cache.</p> <p>Note Age resets each time an entity is accessed by a query engine.</p> <p>Default is 900.</p>

Table 56 CDD Agent Classes Tab Inference Agent and Query Agent Settings (Sheet 2 of 3)

Setting	Notes
Queue Size (Shared Queue)	<p>Used for destinations whose threading model is Shared Queue (see Threading Model in Table 55, CDD Collections Tab Input Destination Settings, on page 494).</p> <p>Specifies the queue size for the processing unit-wide shared queue</p> <p>Note: In this release, set the same value for all agents configured to deploy in the same processing unit.</p> <p>If set to 0 (zero), the queue size is unlimited.</p> <p>Default is 1024</p>
Thread Count (Shared Queue)	<p>Used for destinations whose threading model is Shared Queue (see Threading Model in Table 55, CDD Collections Tab Input Destination Settings, on page 494).</p> <p>Specifies the number of processing unit-wide shared threads.</p> <p>Note: In this release, set the same value for all agents configured to deploy in the same processing unit.</p> <p>As a guideline, set the value to the number of processors available to the JVM.</p> <p>In MM Console, this thread appears with the name <code>\$default.be.mt\$</code>.</p> <p>Default value is 10.</p>
Max Active	<p>Specifies the maximum number of active agents of this class. This value is used for fault tolerance. Deployed agents that are acting as standbys can take over from active instances that fail. In many cases, there is no need to keep standby instances.</p> <p>A value of 0 indicates an unlimited number of active instances.</p> <p>See Fault Tolerance of Agents in <i>TIBCO BusinessEvents Architect's Guide</i> for more details.</p> <p>Default is 0.</p>

Table 56 CDD Agent Classes Tab Inference Agent and Query Agent Settings (Sheet 3 of 3)

Setting	Notes
Inference Agent Settings	
BusinessWorks Repo URL	<p>If this project will integrate with a TIBCO ActiveMatrix BusinessWorks project, enter the Repo URL for the ActiveMatrix BusinessWorks project repo URL here.</p> <p>Use forward slashes.</p> <p>See Overview of Integration with ActiveMatrix BusinessWorks on page 638 for more details.</p>
Concurrent RTC	<p>If checked, enables concurrent run to completion cycles, generally shortened to RTC cycles. (Also known in prior releases as concurrent Rete and concurrentwm).</p> <p>Concurrent RTC does not require cache OM but does require local locking.</p> <p>The number of concurrent cycles is determined by the number of available threads. See Chapter 29, Agent and Processing Unit Configuration for details. Also see Chapter 9, Concurrency and Project Design in <i>TIBCO BusinessEvents Architect's Guide</i> for important information on using concurrency features.</p>
Check for Duplicates	<p>By default, TIBCO BusinessEvents checks if the external IDs (@extId) of entities are unique within the current agent. If you want to check for uniqueness of external IDs across the cluster, check this check box. Performing this check affects performance.</p> <p>Default is unchecked.</p>

CDD Agent Classes Tab Properties Reference

Properties are available for inference agents, cache agents, and query agents.

Table 57 CDD Agent Classes Tab Inference Agent and Query Agent Properties (Sheet 1 of 5)

Setting	Notes
Inference Agent and Query Agent Properties	
<code>com.tibco.cep.runtime.channel.payload.validation</code>	<p>XML event payloads are validated when this property is set to true. There may be some loss of performance due to the extra processing.</p> <p>Default is false.</p>
Inference Agent Properties	
<code>Agent.AgentClassName.recoveryPageSize</code>	<p>Specifies the number of entries per page to be used while recovering objects from the cache.</p> <p>For example, if you set the value to 10,000, then the engine loads handles in blocks of 10,000, instead of trying to load them in a single batch. Smaller batch sizes result in slower recovery. Experiment with batch size to establish the best batch size to use for your environment.</p> <p>A value of 0 means that the objects are recovered in one iteration.</p> <p>Default is 0.</p>

Table 57 CDD Agent Classes Tab Inference Agent and Query Agent Properties (Sheet 2 of 5)

Setting	Notes
<code>Agent.AgentClassName.cacheTxn.updateCache</code>	<p>Used only if cache-aside database write strategy is used.</p> <p>If set to false: When a rule action changes the value of an entity's properties, then the entity instance is evicted from the cache instead of updating it. Updates are saved in the backing store. Use this setting and <code>Agent.AgentClassName.threadcount</code> as needed to improve performance and cache memory management.</p> <p>This property interacts with the Cluster > Domain Objects setting, Evict From Cache on Update (and its override settings if any):</p> <ul style="list-style-type: none"> • When this CDD property is set to true, the domain objects Evict From Cache on Update setting is ignored, in the agent for which the property is set. • When this CDD property is set to false, the domain objects Evict From Cache on Update setting overrides this CDD property. <p>See CDD Cluster Tab Domain Objects Default Settings Reference on page 481 for details on the Evict From Cache on Update setting.</p> <p>Possible values are true and false.</p> <p>Default is true.</p>
<code>Agent.AgentClassName.threadcount</code>	<p>For use with cache aside and only when parallel operations feature is used (see <code>Agent.agentClassName.enableParallelOps</code>).</p> <p>Defines the number of <code>\$CacheWriter</code> threads performing cache writing jobs.</p> <p>See Chapter 10, Threading Models and Tuning in <i>TIBCO BusinessEvents Architect's Guide</i> for usage guidelines.</p> <p>Tip This property is also used to define the number of Recovery threads (used for recovering Cache Plus Memory entity handles at inference engine startup).</p> <p>Default value is 2.</p>

Table 57 CDD Agent Classes Tab Inference Agent and Query Agent Properties (Sheet 3 of 5)

Setting	Notes
<code>Agent.agentClassName.checkDuplicates</code>	<p>Affects how TIBCO BusinessEvents checks uniqueness of entity external IDs (@extId).</p> <p>If set to false, checks for uniqueness of external IDs within the agent</p> <p>If set to true, checks for uniqueness of external IDs across the cluster. Performing this check affects performance so use it with care.</p> <p>Default is false.</p>
<h3>Inference Agent Database Writer Thread Tuning Properties</h3>	
<p>Note For use with cache aside and only when the parallel operations feature is used.</p> <p>For a guide to usage, see the section Database Write Tuning Options for Cache Aside in Chapter 10, Threading Models and Tuning in <i>TIBCO BusinessEvents Architect's Guide</i>.</p>	
<code>Agent.agentClassName.enableParallelOps</code>	<p>If true, parallel operations are used Post-RTC phase operations are done in parallel:</p> <ul style="list-style-type: none"> • Writes to the cache • Writes to the database (relevant only cache aside strategy is used) • Executes the actions list, for example, sends messages (events) and acknowledges events, as needed. <p>Use of parallel operations generally requires use of locking to ensure data integrity.</p> <p>If false, sequential operations are used All post-RTC phase operations are done on a single thread in the order shown above.</p> <p>This property is set to false for specific needs such as when Caller's Thread threading option is used. For an example, see Using CLIENT_ACKNOWLEDGE Mode with Websphere MQ and Cache-Aside on page 91.</p> <p>Another reason to set the value to false is to ensure that the system waits to send a reply event confirming that some work has been done, until the result of the work can be seen in the cache.</p> <p>Defaults to true only if cache-aside write strategy and concurrent RTC are both used. Otherwise defaults to false.</p>

Table 57 CDD Agent Classes Tab Inference Agent and Query Agent Properties (Sheet 4 of 5)

Setting	Notes
<code>Agent.AgentClassName.dbthreadcount</code>	<p>Defines the number of database write threads available to process the RTC transactions from the queue, that is, the number of threads performing database writing jobs (\$DBWRITER thread pool). Writes include applying entity inserts, updates, and deletes to the database.</p> <p>Although the limit is seldom reached, you can guarantee that a connection is always available for a dbwriter thread as follows. Set this field to the same value as CDD Cluster tab > Backing Store > Connection > Max Size field.</p> <p>Default is 2.</p>
<code>Agent.AgentClassName.dbOpsQueueSize</code>	<p>The size of the queue (a Java blocking queue) for database writing jobs.</p> <p>Zero (0) or a negative value means the queue size is unlimited.</p> <p>Note When the queue is full, all engine operations are blocked.</p> <p>Default is 8.</p>
<code>Agent.AgentClassName.dbOpsBatchSize</code>	<p>Used in the post-RTC phase. Sets the maximum number of RTC transactions that a database writer thread takes from the database operations queue and processes in one batch.</p> <p>Database write threads process the RTC transactions from the queue. The number of threads is defined by <code>dbthreadcount</code>.</p> <p>A database write thread takes <i>up to</i> the <code>dbOpsBatchSize</code> number of RTC transactions, processes them and commits them to the database. (When database write threads are idle, they take available jobs from the database operations queue, even if there are less jobs than <code>dbOpsBatchSize</code>.)</p> <p>Default is 10.</p>

Table 57 CDD Agent Classes Tab Inference Agent and Query Agent Properties (Sheet 5 of 5)

Setting	Notes
Query Agent Properties	
<code>be.agent.query.localcache.prefetchaggressive</code>	<p>If set to true, then objects required for a query are prefetched while the query is executing.</p> <p>The prefetch feature improves performance, but CPU and memory usage increases as a result of the aggressive prefetching. You may have to try different values till you find the optimal settings for your environment.</p> <p>Ensure that the cache size is large enough to accommodate objects that are prefetched.</p> <p>Default is <code>false</code>.</p>
<code>be.agent.query.enable.filter.optimizer</code>	<p>Used only with Oracle Coherence as the cache provider.</p> <p>If set to true, the query agent attempts to use indexing that is enabled and defined as explained in the section Indexing for More Efficient Cache Queries on page 311.</p> <p>If set to false, indexes are ignored by this agent.</p> <p>See also Present in Index, in CDD Cluster Tab Domain Objects Default Settings Reference on page 481.</p> <p>Default is <code>false</code>.</p>
Cache Agent Properties	
<code>be.engine.cacheServer.channel.disable</code>	<p>By default cache agents connect to channels. In most cases, however, cache agents do not need to connect to channels. To prevent cache agents from connecting to channels, set this property to true.</p> <p>Default is <code>false</code>.</p>

Understanding and Configuring Log Configurations

Each processing unit references a log configuration. The log configurations are defined in the Collections tab. See [Configuring Collections of Rules, Rule Functions, and Destinations on page 492](#) for related procedures. This section explains more about log configurations and how to configure them.

You can also replace the default line layout implementation with your own.

Log File Settings

For a reference to the settings, see the [Files Section](#) in [Table 58, CDD Collections Tab Log Configurations Settings](#), on page 509.

Log File Name and Location

Set the name and location of the log file for a log configuration using the Name and Directory settings. If you do not enter a leading slash, the files are stored relative to the working directory (the directory in which you start the `be-engine.exe` executable). If you do not specify a name, the engine name is used. If no engine name is specified the name defaults to `cep-engine.log`.

Number and Size of Log files

You can also set the size of a single log file, the number of files to keep, and whether a log file is flushed when an engine starts, or whether entries are appended.

Understanding Log Configuration Levels and Syntax

In a log configuration that uses the provided line layout implementation, you select a *level* of logging for each *module* in the TIBCO BusinessEvents runtime.

Levels

A level corresponds to how much logging is filtered out. They are ordered where `all` is lowest and `off` is highest:

Level	Description
Off	Highest possible rank. Filters out all logging messages (turns logging off for the specified module).

Level	Description
Fatal	Logs only severe runtime errors that cause the application to stop running.
Error	Also logs runtime errors that might not cause the application to stop running.
Warn	Also logs potentially harmful runtime events or situations.
Info	Also logs runtime informational events of general interest.
Debug	Also logs detailed runtime informational events, for use in identifying issues.
Trace	Also logs even more detailed runtime information.
All	Lowest possible rank. Turns on all logging including any custom logging levels.

Syntax

Enabling a lower level automatically enables the higher levels. For example, enabling `info` automatically enables `fatal`, `error`, and `warning`.

Assign each module to a level using a space-delimited list. The levels are not case sensitive. The syntax is as follows:

```
module1:level module2:level . . .
```

To assign a certain level of logging to *all* modules, use an asterisk:

```
*:info
```

This syntax means that logging for all modules is at the `info` logging level.

You can use the asterisk syntax and also specify exceptions that use a different logging level. For example:

```
*:info driver.tibrv:debug
```

This syntax means that all modules use logging level `info`, except the module `driver.tibrv` which uses `debug` level.

Configuring Log Configurations

See [Understanding and Configuring Log Configurations on page 506](#) for an explanation of the logging levels, modules, and syntax details

To Add a Log Configuration

1. In the Collections tab select Log Configurations and click **Add**.

2. In the Configuration section, give the log configuration a name.
3. Add the log levels you want to enable in this configuration. See [CDD Collections Tab Log Configurations Settings Reference on page 509](#) for details on the fields.
4. If you want to redirect the STDERR or STDOUT streams, check the Enable checkbox and follow guidelines in the [Send to Terminal Section on page 510](#).
5. Save.

To Specify a Custom Line Layout Class

1. In the Collections tab select Log Configurations and click **Add**.
2. In the Configuration section, give the log configuration a name.
3. In the Custom Line Layout section, check the Enable checkbox, and complete the Class and Argument fields as shown in [Table 58, CDD Collections Tab Log Configurations Settings, on page 509](#).
4. Save.

CDD Collections Tab Log Configurations Settings Reference

Additional log settings can be added to the processing Unit tab Properties. See [CDD Processing Units Tab Settings Reference on page 515](#).

Table 58 CDD Collections Tab Log Configurations Settings (Sheet 1 of 3)

Property	Notes
Name	Name of this log configuration.
Enable	Check the Enable checkbox to enable this log configuration. All other Enable settings are ignored if this checkbox is unchecked.
Levels	Space-separated list of levels and modules used in this log configuration. See Syntax on page 507 and other sections in introduction to Understanding and Configuring Log Configurations on page 506 Default is info
Files Section	
Enable	Check the Enable checkbox to enable log files to be written. Configure the settings in this section to specify details. If this checkbox is unchecked, all other properties in this section are ignored.
Directory	Enter the absolute path to the directory in which you want to store the files. If you do not enter a leading slash, the files are stored relative to the working directory (the directory in which you start the <code>be-engine.exe</code> executable).
Name	Name of the log file. The default value is the engine name. If no engine name is set, then the default value is <code>cep-engine.log</code>

Table 58 CDD Collections Tab Log Configurations Settings (Sheet 2 of 3)

Property	Notes
Max number	<p>Number of log files to keep. When the Max size setting value is reached, a new log file is created for the next log entries. Files are created up to the Max number setting size. The oldest file is deleted when a new file is added after this value is reached.</p> <p>Default is 10.</p>
Max size	<p>Maximum size of one log file.</p> <p>Default is 10000000.</p>
Append	<p>If checked then new entries are added to the end of the file. If not checked, the contents of the file are flushed each time the engine starts.</p>
Send to Terminal Section	
Enable	<p>Check the Enable checkbox to enable the redirections specified in this section. If this checkbox is unchecked, all other properties in this section are ignored.</p>
Output redirection	<p>If true, the STDOUT stream is written to the terminal. If false, it is not.</p>
Error redirection	<p>If true, the STDERR stream is written to the terminal. If false, it is not.</p>
Custom Line Layout Section	
Enable	<p>Check the Enable checkbox to enable the custom line layout entries to take effect. Configure the settings in this section to specify details of a custom layout.</p> <p>If this checkbox is unchecked, all other properties in this section are ignored.</p> <p>If this checkbox is checked all properties in the other sections are ignored (except Name, and Enable in the upper section).</p>

Table 58 CDD Collections Tab Log Configurations Settings (Sheet 3 of 3)

Property	Notes
Class	<p>The custom line layout class.</p> <p>This class must implement <code>org.apache.log4j.Layout</code> and must be available in the runtime classpath.</p> <p>The class needs 2 constructors:</p> <ul style="list-style-type: none">• One with no argument• One with a single String argument, which receives the value of the Arguments field.
Arguments	<p>A String parameter used for the custom line layout class, if required:</p> <ul style="list-style-type: none">• To use the constructor that requires an argument, specify the argument• To use the constructor that does not expect an argument, leave the field empty.

Logging for TIBCO BusinessEvents DataGrid

You must configure logging for TIBCO BusinessEvents DataGrid in TIBCO BusinessEvents separately. To do so, set the following properties in the CDD file:

Property	Description
be.engine.cluster.as.log.dir	The directory to which the TIBCO BusinessEvents DataGrid log files will be written. If unspecified, the logs will be written to the same directory as the TIBCO BusinessEvents logs.
be.engine.cluster.as.log.filename	File name of the log file. By default, the file name is as- <i><pid></i> .log
be.engine.cluster.as.log.level	The log level specifying how much logging is to be filtered out. See Log Configuration Levels for TIBCO BusinessEvents DataGrid Logging for details. By default, the log level is set to INFO.

Log Configuration Levels for TIBCO BusinessEvents DataGrid Logging

Following levels can be used to specify the level of logging for TIBCO BusinessEvents DataGrid logs. Note that the logging levels are case insensitive.

Level	Description
None	Highest possible rank. Filters out all logging messages (turns logging off for the specified module).
Fatal	Logs only severe runtime errors that cause the application to stop running.
Error	Also logs runtime errors that might not cause the application to stop running.
Warn	Also logs potentially harmful runtime events or situations.
Info	Also logs runtime informational events of general interest.
Fine	Also logs detailed runtime informational events, for use in identifying issues.
Finer	Also logs even more detailed runtime information.
Finest	Lowest possible rank. Turns on all logging including any custom logging levels.

Configuring Processing Units (All OM Types)

To configure a processing unit (PU), you add the items you configured earlier, and any additional properties required. If you do not find a configuration item you require, click the appropriate tab and add it, then return to Processing Units tab and continue configuration.

One processing unit named default is provided out of the box. You can change this name. It has no significance, except that TIBCO Administrator expects a processing unit of this name by default, which can be useful for testing purposes.

Processing units are referenced in the site topology file used by TIBCO BusinessEvents Monitoring and Management component (see Chapter 3, Basic MM Configuration in *TIBCO BusinessEvents Administration*).

PU with Unique Agent Instance Properties

Depending on configuration, some processing units can be deployed more than once in a cluster. Others have unique configuration details that make them deployment-specific, that is, that limit them to being deployed only once in a cluster. Configuring the following CDD settings for a processing unit makes it deployment-specific processing unit:

Agent Key At the processing unit node, you can associate a unique key with a selected agent class. This key identifies an agent instance uniquely at runtime. The purpose of the agent key is to retrieve scorecards from the backing store. Scorecards are local to an agent and the key enables the correct scorecard to be returned to the correct agent.

Agent Priority The agent priority determines which agents of a given class are active, when fault tolerance is used. Each deployed agent of an agent class must have a different priority.

See Deployment-Specific Processing Units in *TIBCO BusinessEvents Administration* for other ways a processing unit can be deployment-specific.

To Add a Processing Unit

See [CDD Processing Units Tab Settings Reference on page 515](#) for guidelines on the settings and properties.

1. At the Processing Units tab do the following:
 - Select the default processing unit and configure it. You can rename it as needed.
 - Click **Add** to add more processing units as needed.

2. In the Name field, enter the name for the processing unit as needed.
For deployment, TIBCO Administrator by default looks for a processing unit called default and a CDD file called default.
3. In the Log Configuration field, browse to and select one log configuration. See [Understanding and Configuring Log Configurations on page 506](#) for more details on log configurations.
4. Check the Hot Deploy checkbox if you want to enable hot deployment. See *TIBCO BusinessEvents Administration* for details about hot deployment.
5. If you use the TIBCO BusinessEvents Data Modeling add-on, check the Enable DB Concepts checkbox to enable database concept functionality on this processing unit as desired.
6. In the Agents section, click **Add** and select an agent class.
7. If needed, assign to each agent a key and a priority. See [PUs with Unique Agent Instance Properties on page 513](#) for details.
8. In the Properties section, add any additional configuration properties as required. For example see [Localhost and Localport Properties on page 439](#) for one use case.

CDD Processing Units Tab Settings Reference

Table 59 CDD Processing Units Tab Settings (Sheet 1 of 2)

Property	Notes
Name	Enter a name that is unique across the cluster.
Log Configuration	Browse to and select a log configuration, configured at Collections tab. See Understanding and Configuring Log Configurations on page 506 for more details.
Hot Deploy	Check the checkbox to enable hot deployment for this processing unit. See <i>TIBCO BusinessEvents Administration</i> for details about hot deployment.
Enable Cache Storage	<p>Check the checkbox to enable cache storage on this processing unit (PU). Settings available depend on the types of cache agents in the PU.</p> <ul style="list-style-type: none"> • In PUs used to host cache agents: The checkbox is checked and cannot be unchecked. • In PUs used only for dashboard agents (available in TIBCO BusinessEvents Views only): the checkbox is unchecked and cannot be checked. Dashboard agents cannot function as cache servers. • in PUs that host inference agents or query agents (or both): <ul style="list-style-type: none"> — If checked, the PU is used for storing cache data. — If unchecked, the PU is not used for storing cache data. <p>Note: Enable cache storage in PUs running inference and query agents for test deployments only. Not recommended in production.</p> <p>Default value for PUs containing inference agents or query agents (or both) is unchecked.</p>
Enable DB Concepts	Check the checkbox to enable database concepts functionality for this processing unit. Available only with TIBCO BusinessEvents Data Modeling add-on software.

Table 59 CDD Processing Units Tab Settings (Sheet 2 of 2)

Property	Notes
Agents Section	
Agent	<p>Name of the agent class you selected. Agent classes are defined at the Agent Classes tab.</p>
Key	<p>Specifies a value that uniquely identifies an instance of an agent of this class at deploy time.</p> <p>Required for recovery of scorecards. Recommended in all cases, for situations that require an agent instance to be uniquely identified.</p> <p>The value for Key must uniquely identify the agent.</p> <p>Note: In certain TIBCO BusinessEvents Monitoring and Management methods, you may be prompted for a session name. For session name you generally put the agent class name. However, if the agent class also has a key, you must instead use the key value. Because of this, the key value must uniquely identify the agent.</p> <p>No default value.</p>
Priority	<p>Specifies the priority of the agent for load balancing purposes. (Not available in TIBCO BusinessEvents Express.)</p> <p>The priority indicates the order in which standby agents become active, and conversely, the order in which active agents become standbys, when new agents join the cluster.</p> <p>The <i>lower</i> the number, the higher the agent is in the activation priority list. For example, an agent with priority 2 has a higher priority than an agent with a priority of 6.</p> <p>Ensure that inference agents of the same class in different PUs have different values. This value determines the order of each instance of an agent class for startup, as well as failover and failback in fault tolerance situations.</p> <p>No default value.</p>

CDD Processing Units Tab Coherence Log Properties Reference

The properties in the Coherence Log Properties section are used to configure the Coherence log. This log is used only if Oracle Coherence is used as the cache provider.

Standard logging settings are configured in the Log Configuration tab. See [CDD Collections Tab Log Configurations Settings Reference on page 509](#).

Table 60 CDD Processing Units Tab Coherence Properties (Sheet 1 of 3)

Property	Notes
Coherence Log Properties	
<code>tangosol.coherence.log</code>	<p>Specifies the output device used by the logging system.</p> <p>Optional.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • <code>stdout</code> • <code>stderr</code> • <code>jdk</code> (Requires JDK 1.4 or later) • <code>log4j</code> (Requires log4j libraries to be in the classpath) • A file path and file name <p>If you specify <code>jdk</code> or <code>log4j</code> you must also perform appropriate configuration of the JDK or Apache <code>log4J</code> logging libraries.</p> <p>Default is <code>stdout</code>.</p>

Table 60 CDD Processing Units Tab Coherence Properties (Sheet 2 of 3)

Property	Notes
<code>tangosol.coherence.log.level</code>	<p>Specifies which logged messages are output to the log destination.</p> <p>Optional.</p> <p>Possible values are:</p> <ul style="list-style-type: none">• 0: Only output without a logging severity level specified will be logged• 1: All the above plus errors• 2: All the above plus warnings• 3: All the above plus informational messages• 4-9: All the above plus internal debugging messages (the higher the number, the more the messages)• -1: No messages <p>Default is 5.</p>
<code>tangosol.coherence.log.limit</code>	<p>Specifies the maximum number of characters that the logger daemon processes from the message queue before discarding all remaining messages in the queue.</p> <p>The message that causes the total number of characters to exceed the maximum is not truncated.</p> <p>All discarded messages are summarized by the logging system with a single log entry detailing the number of discarded messages and their total size. When the queue empties, the logger is reset and subsequent messages are again logged.</p> <p>The purpose of this setting is to avoid a situation where logging can itself prevent recovery from a failing condition, for example by contributing to timing issues.</p> <p>Logging occurs on a dedicated low-priority thread to further reduce its impact on the critical portions of the system.</p> <p>Optional.</p> <p>Possible values are positive integers or zero (0). Zero implies no limit.</p> <p>Default is 0</p>

Table 60 CDD Processing Units Tab Coherence Properties (Sheet 3 of 3)

Property	Notes
JMX Management properties	
You must specify the following JMX-related properties to enable cluster statistics to appear in the monitored objects table in MM. (MM is not available in TIBCO BusinessEvents Express.)	
<code>tangosol.coherence.management</code>	Default is all
<code>tangosol.coherence.management.remote</code>	Default is true

CDD Processing Units Tab JMS Server Connection Properties

The following properties relate to JMS channels. Most enable you to define how TIBCO BusinessEvents attempts to reconnect to a JMS server in the event of a disconnection. Two relate to disabling connections. See [Chapter 5, JMS Channels](#), on [page 83](#) for details about configuring a JMS channel.

You can add these properties at the cluster level if they apply to JMS channels in all processing units in the cluster.

Table 61 CDD Processing Units Tab Properties for Reconnecting to a JMS Server

Property	Notes
<code>com.tibco.tibjms.connect.attempts</code>	<p>Specifies the number of reconnection attempts, and the interval between each attempt to connect to the JMS server.</p> <p>The value must use the format: <i>attempts, retry interval</i>. For example: 10, 500 means 10 attempts, with a 500 millisecond interval between each retry attempt.</p> <p>This property is used only for channels that have a TIBCO Enterprise Message Service provider.</p> <p>Note: Use <i>either</i> <code>be.jms.reconnect.timeout</code> or <code>com.tibco.tibjms.connect.attempts</code>. If you set both the properties, then <code>com.tibco.tibjms.connect.attempts</code> takes precedence.</p> <p>Default is 2, 500</p>
<code>be.jms.reconnect.timeout</code>	<p>Specifies the retry interval (in seconds) for reconnecting to the JMS server when the connection is broken.</p> <p>A value of zero (0) means do not retry. Any other value means keep retrying (with no limit to number of retries), and use the specified interval between each attempt.</p> <p>Note: Unacknowledged messages (Events) are resent to the TIBCO BusinessEvents engine, which may result in duplicate events.</p> <p>Note: Use <i>either</i> <code>be.jms.reconnect.timeout</code> or <code>com.tibco.tibjms.connect.attempts</code>. If you set both the properties, then <code>com.tibco.tibjms.connect.attempts</code> takes precedence.</p> <p>Default is 0 (zero)</p>

Table 61 CDD Processing Units Tab Properties for Reconnecting to a JMS Server (Cont'd)

Property	Notes
<code>be.jms.reconnect.msgCodes</code>	<p data-bbox="392 291 1316 355">Specifies a case-insensitive character pattern that matches all error messages or error codes that will cause a reconnect attempt.</p> <p data-bbox="392 373 1245 437">This property is used for JMS channels with providers other than TIBCO Enterprise Message Service.</p> <p data-bbox="392 454 1109 484">Default is * (that is, the wildcard matched by any characters.)</p>
<code>be.channel.tibjms.queue.disabled</code>	
<code>be.channel.tibjms.topic.disabled</code>	<p data-bbox="392 604 1316 703">By default, be-engine connects to all defined channels on startup, including those not mentioned in the CDD file. This is because such channels can be used as output channels. However this is not always desired.</p> <p data-bbox="392 720 1306 847">To disable queue or topic connections for specific JMS channels, add the following properties as appropriate. Enter the project path to the JMS channel as the individual value. Use commas or spaces as the delimiter. Use forward slashes in the project path. For example:</p> <pre data-bbox="392 864 1168 916">be.channel.tibjms.queue.disabled=channels/1jmschannel, channels/3jmschannel</pre> <pre data-bbox="392 933 1168 986">be.channel.tibjms.topic.disabled=channels/2jmschannel channels/4jmschannel</pre>

Chapter 30 JDBC Backing Store Setup

A backing store enables persistent backup of the objects generated and modified at runtime. Use of a backing store enables recovery in the event of a system-wide failure.

For instructions on migrating from the legacy Oracle Types (Oracle only) backing store to the JDBC backing store, see the migration chapters in *TIBCO BusinessEvents Installation*.



TIBCO BusinessEvents Express This chapter relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

Topics

- [JDBC Backing Store Setup and Configuration Overview, page 524](#)
- [Cases That May Need Additional Setup, page 527](#)
- [Resources Required for Setting Up the Database, page 529](#)
- [DBMS Software Requirements and Installation, page 531](#)
- [Configure the CDD for Special Cases \(As Needed\), page 536](#)
- [Add a JDBC Connection Resource \(Now or Later\), page 537](#)
- [Configure Backing Store Settings in the CDD \(Now or Later\), page 539](#)
- [Build the EAR File, page 542](#)
- [Initialize the Database and Generate Non-Project Tables, page 543](#)
- [Generate the Project-Schema-Specific SQL Scripts, page 545](#)
- [Check the Aliases File, Run the Project Schema Script, page 549](#)
- [Updating an Existing Backing Store Schema](#)
- [Handling Entities Deleted from a Write-Behind Backing Store, page 555](#)
- [Backing Store Table Reference, page 556](#)

JDBC Backing Store Setup and Configuration Overview

The backing store feature requires use of Cache object management. Before you add a backing store, develop your caching solution and test it. Also ensure that your project ontology is completely configured.

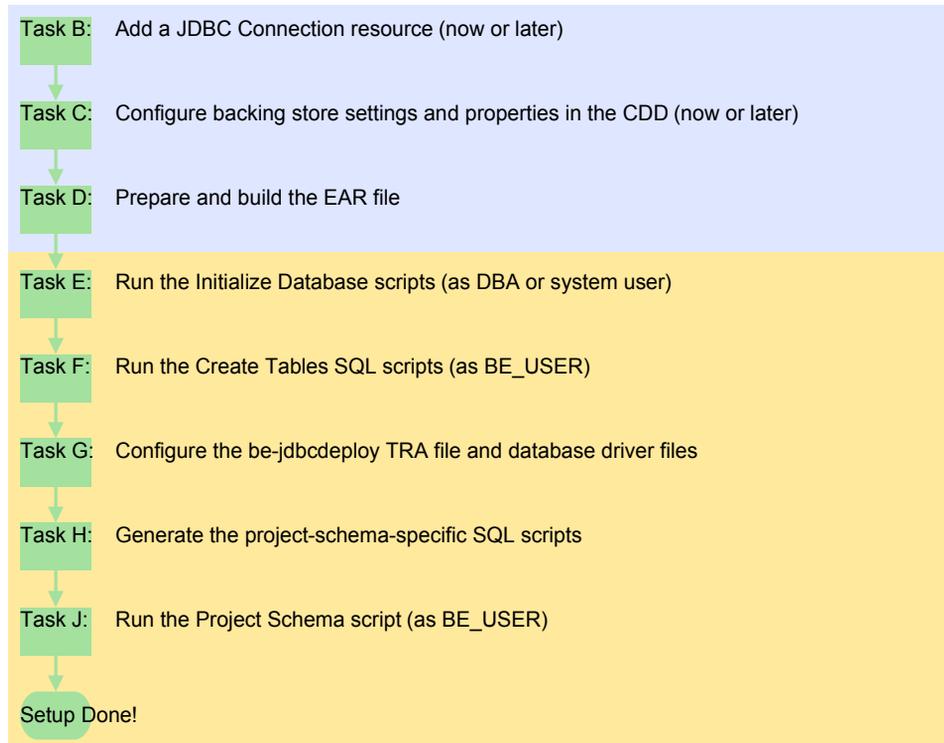


- Use a separate schema and schema owner for each project, even if different projects use the same ontology (otherwise ontology conflicts can occur).
- If your project ontology changes after the backing store is in place, you must update the backing store schema. See [Updating an Existing Backing Store Schema on page 552](#).

The upper (blue) area in the following diagrams shows TIBCO BusinessEvents Studio configuration tasks. The lower (yellow) area shows database setup utility tasks. The tasks shown map to task sections in this chapter.

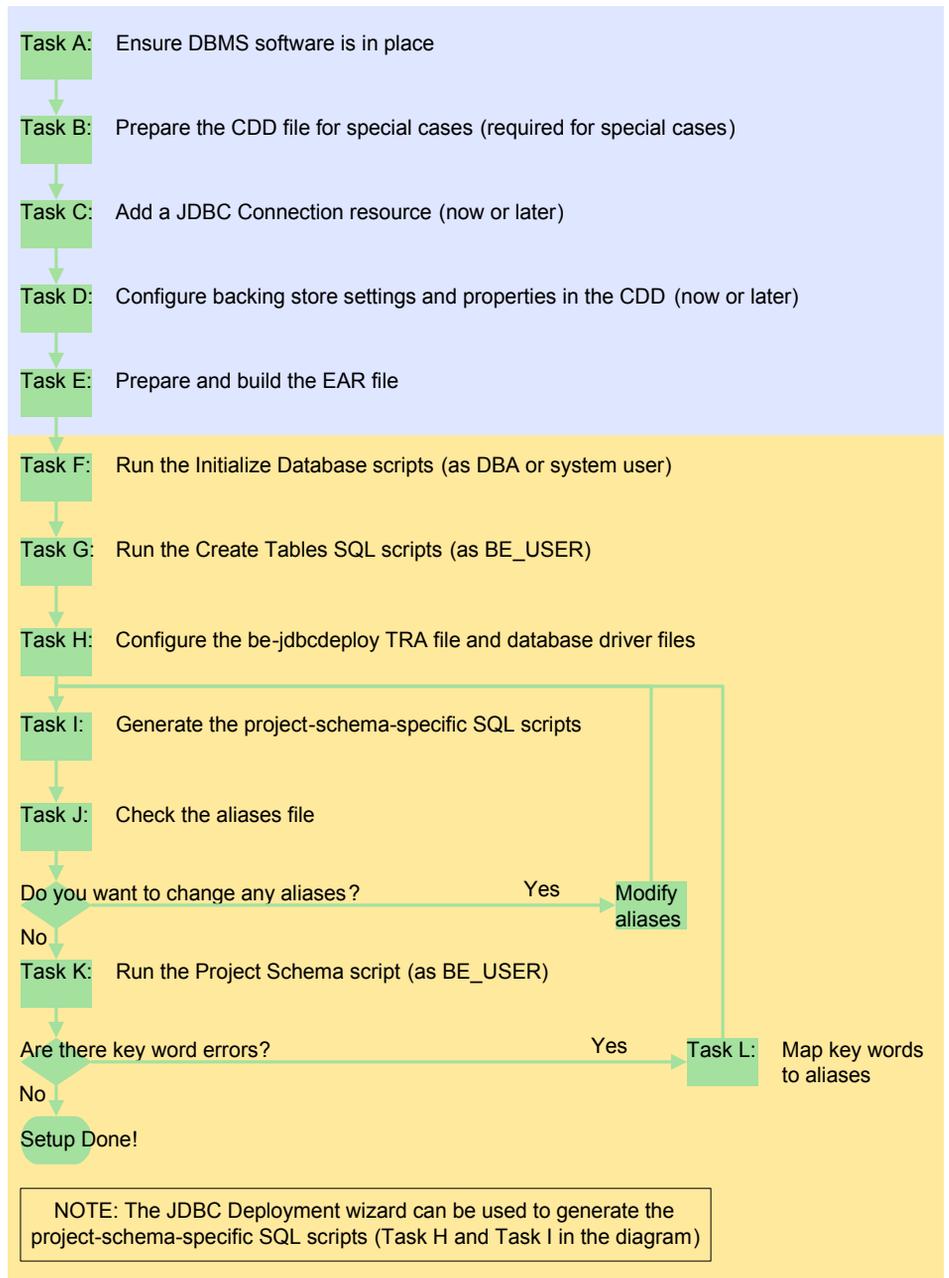
Simplest Case

For the simplest case where no additional project configuration is required (see [Cases That May Need Additional Setup on page 527](#)), the main tasks in setting up a backing store are as shown next.



Most Complex Case

The complete flow of tasks for all possible situations is shown here. For more information see [Cases That May Need Additional Setup on page 527](#).



Backing Store Setup

Setup refers to using the provided scripts to create the backing store schema for your project. See [Resources Required for Setting Up the Database on page 529](#) for the DBMS-related requirements.

For the basic setup tasks, see [Initialize the Database and Generate Non-Project Tables on page 543](#). However it is important that you first read the section [Cases That May Need Additional Setup on page 527](#) to understand if you need to do any special project configuration before running the scripts. (Other project configuration can be done before or after you do the setup tasks.)

Backing Store Configuration

In addition to setting up the backing store (and doing related project configuration for special cases), you also configure the TIBCO BusinessEvents Studio project to use the backing store according to your needs. You can do the following to configure backing store behavior:

- Use the Shared All persistence option, with Database Type Oracle or SQL Server. (Formerly known as JDBC backing store.)
- You can use either write-behind or cache-aside database write strategy.
- Tune the database connection pool properties.
- You can use Use a limited or unlimited cache. You can define a global setting and configure object level overrides.
- Control how the cache is preloaded from the backing store at startup. You can define a global setting and configure object level overrides.

See [Configure Backing Store Settings in the CDD \(Now or Later\) on page 539](#).

You must also add a JDBC Connection resource to your project, before or after the backing store setup.

To make the flow of tasks simpler, all TIBCO BusinessEvents Studio project configuration is documented before the setup tasks, because some TIBCO BusinessEvents Studio configuration is required in certain cases. The procedures make it clear which configuration can be done before or after backing store setup.

Cases That May Need Additional Setup

This section explains some cases that may require additional setup.

Ontology Identifiers that Exceed the DBMS Maximum Column Length

Entity names and entity property names are used by backing store scripts to generate database table and column identifiers.

DBMSs put different limits on the length of a database identifier name. For example, in Oracle the maximum length is 30 characters, and in SQL Server the limit is 128 characters.

Generated database identifiers are longer than the TIBCO BusinessEvents identifiers because they contain characters in addition to the TIBCO BusinessEvents identifier.

You can handle long identifier issues in either of the following ways (or a combination of these ways).

Letting the Utility Generate Short Aliases for Table Names

You can allow the `be-jdbcdeploy` utility to generate short aliases for long names. You can also edit those names and rerun the utility. For details see [Task I, Check the Aliases File and Modify Aliases as Desired, on page 549](#).

Note that alias file entries are also generated for another reason. See [Ontology Identifiers that Use Database Key Words on page 528](#).

Specifying Short Table Names in the CDD File

You can avoid the problem of long entity type names before you begin to configure the backing store by specifying short database identifiers using the CDD Table Name setting.

The advantage of this method is that you can choose meaningful names before running the `be-jdbcdeploy` utility. The disadvantages are that you may not know ahead of time which entities require short names, and you must also ensure that the table names you specify are unique across all entities in the ontology.

If you do not specify table names, and entity names are repeated, on the other hand, the generated table names are appended with dollar (\$) characters as necessary, for example, `D_ORDER`, `D_ORDER$`, `D_ORDER$$` and so on.

See [Configure the CDD for Special Cases \(As Needed\) on page 536](#) for details.

Ontology Identifiers that Use Database Key Words

As well as database names that are too long, ontology terms that are key (reserved) words in your DBMS product must also be mapped to an alias. If errors occur when you run the SQL scripts due to key word clashes, examine the errors and add the appropriate words to the key word mapping file.

A provided file (*BE_HOME/bin/dbkeywordmap.xml*) ships with some basic mappings: *start*, *end*, *schema*, *mode*, and *index*. You can use it as a model.

Unlike the Aliases file, the key word mapping file is not a project-specific file. It is intended to be generally useful across different projects. However, keyword mappings are also added to the aliases file when you run the SQL scripts, so you can also provide project-specific aliases for the generic mappings, if you want to.

The procedures are explained in the task sections within the section [Initialize the Database and Generate Non-Project Tables on page 543](#).

String Properties that Exceed the DBMS Maximum Column Length

The default column size for String type attributes is 255 characters. If you expect the data length of an entity property to exceed that value, then in the CDD file set the Max Length field for each entity's properties. The utility changes the data type of String attributes with long lengths to CLOB, as appropriate.

See [Configure the CDD for Special Cases \(As Needed\) on page 536](#) for details.

Excluding Entities from the Backing Store

You do not have to use the backing store for all entities. In the CDD file you can specify entities for which you do not want to use the backing store.



If later you want to include any excluded entities, you must change the setting and update the backing store setup as explained in [Updating an Existing Backing Store Schema on page 552](#).

See [Configure the CDD for Special Cases \(As Needed\) on page 536](#) for details.

Resources Required for Setting Up the Database



For specific database products and versions supported, see the readme file, which is in the `TIBCO_HOME/release_notes` directory.

Provided Configuration Resources

The table below lists resources required and sections following explain the procedures for setting up backing store tables.

Table 62 Resources Required for JDBC Backing Store Implementation

Resource	Default Location and Notes
Provided Files in <code>BE_HOME/bin</code>	
<code>base_types.xml</code>	The <code>base_types.xml</code> file is used by the deployment utility. Do not edit this file.
<code>be-jdbcdeploy executable</code> <code>be-jdbcdeploy.tra</code>	Only used for manual SQL script generation. Generally not needed. You can use a TIBCO BusinessEvents Studio option instead. (See Generated SQL Scripts below).
<code>create_tables_oracle.sql</code> <code>create_tables_sqlserver.sql</code>	Use the appropriate SQL (DDL) script for your DBMS. This script creates the tables that are used to maintain the metadata. The script drops any existing tables and recreates them.
<code>dbkeywordmap.xml</code>	This file contains mappings to handle words used in the TIBCO BusinessEvents project that are database reserved words. See Ontology Identifiers that Use Database Key Words on page 528 for details.

Table 62 Resources Required for JDBC Backing Store Implementation (Cont'd)

Resource	Default Location and Notes
<code>initialize_database_oracle.sql</code>	Use the appropriate script for your DBMS.
<code>initialize_database_sqlserver.sql</code>	<p>By default the user is called <code>BE_USER</code> with the password <code>BE_USER</code> and the user has DBA rights. Edit the script if you want the user to have a different name or different rights.</p> <p>For SQL Server, this script also creates the default database, with the name <code>BE_USER</code> and makes it the default database for the user <code>BE_USER</code>.</p> <p>Note: Use a different user (and schema) for every TIBCO BusinessEvents project that needs a backing store. This script drops the user (and therefore all the tables) and adds the user again.</p>
Generated SQL Scripts	
<p>These scripts are generated when you use the File > Export > JDBC Deployment wizard. The value for <i>yourname</i> is specified in the Output Script Name Prefix setting. You specify the location of the scripts in the wizard</p> <p>You can manually execute the <code>be-jdbcdeploy</code> executable. You specify script name prefix at the command line. Scripts are generated in the same directory where you run <code>be-jdbcdeploy</code>.</p>	
<code>yourname.sql</code>	This SQL (DDL) script creates schema tables and types.
<code>yourname.aliases</code>	This file has entries if the database table identifiers are longer than the DBMS maximum character limit. See and Task I, Check the Aliases File and Modify Aliases as Desired, on page 549 .
<code>yourname_alter.sql</code>	The <code>yourname_alter.sql</code> script is for use in schema migration. Generated only after updates are made to the <code>be-jdbcdeploy.tra</code> file. See Updating an Existing Backing Store Schema on page 552 .
<code>yourname_cleanup.sql</code>	For use as needed. This script truncates the tables.
<code>yourname_delete.sql</code>	This script is used for deleting entities that have been marked as deleted (applies only if write-behind database write strategy is used).
<code>yourname_remove.sql</code>	For use as needed. This script removes the database schema. You can use it to reset the project.

DBMS Software Requirements and Installation

Database Location	Instructions in this chapter assume you are working with a local database for testing purposes. Adapt the instructions if you are working with a remote database. For example, in production environments, you might have to ask a database administrator to create a database user for you. You should then be able to run the other SQL scripts yourself, logged on as the user created by the administrator.
Minimum User Permissions	By default the TIBCO BusinessEvents user permissions are set to DBA privileges. At a minimum, the user must be able to create tables and views. For example for an Oracle database you could use the following:

```
DROP USER BE_USER CASCADE ;
CREATE USER BE_USER IDENTIFIED BY BE_USER;
GRANT CONNECT TO BE_USER;
GRANT RESOURCE TO BE_USER;
GRANT CREATE ANY VIEW TO BE_USER;
GRANT CREATE ANY TABLE TO BE_USER;
```

Task A Ensure DBMS Software and Drivers are in Place

Before you begin to configure the backing store, you or your DBA must do the following:

- Install and start a supported DBMS product. See the product readme file for a list of supported products.
- Copy the appropriate JDBC drivers file to *BE_HOME/lib/ext/tpcl*. You must restart BusinessEvents Studio Explorer after copying the drivers file. This step is required before you can use the design-time Test Connection feature. It is also required for runtime.
- Now or later: If you will use the debugger or tester features, add your DBMS product's libraries to the TIBCO BusinessEvents Studio classpath. See the section [Working with External Library and Custom Function Paths on page 8](#).

The remainder of this section provides a few tips for each supported DBMS.

SQL Server

Here are a few helpful points about SQL Server:

SQL Server authentication for non-production purposes

It is convenient to use SQL Server authentication so you can create database users as needed. Select this option when you install Microsoft SQL Server. With Windows Authentication, on the other hand, you may have difficulties creating users without help from others in your enterprise.

Availability Group

An availability group must be ready and dedicated to TIBCO BusinessEvents so that the TIBCO BusinessEvents database(s) can be added to that group.



It is required to create a full backup of the database before adding it the availability group.

SQL Server AlwaysOn Availability Groups

The AlwaysOn Availability Groups feature in SQL Server 2012 Enterprise Edition is a high-availability and disaster-recovery solution used by TIBCO BusinessEvents. To implement this feature, make sure that your SQL Server setup is up and running properly as per the Microsoft documentation.

Connection Properties to the SQL Server Cluster

Use the availability group listener defined at the SQL Server cluster level in the TIBCO BusinessEvents JDBC resource. The availability group listener enables clients to connect to a SQL Server replica without knowing the name of the physical instance of the SQL Server, for example:

```
jdbc:sqlserver://sqlserver-group-listener:1433;databaseName=be_user
```

To allow faster TCP connection retries, use the property `multiSubnetFailover=true` even if your cluster is on the same subnet as recommended by Microsoft,

Set this property in the BusinessEvents JDBC resource in the URL field, for example:

```
jdbc:sqlserver://sqlserver-group-listener:1433;databaseName=be_user;multiSubnetFailover=true
```

To support Windows authentication, use property `integratedSecurity=true`

See [SQL Server authentication vs. Windows authentication](#) for details.

Windows Authentication

This authentication is only supported on Microsoft Windows operating systems.

In order for TIBCO BusinessEvents to support Windows authentication when accessing SQL Server database, follow these steps:

1. Download and install the Microsoft JDBC 4 Driver distribution on the machines BusinessEvents runs on. The distribution contains the `sqljdbc_auth.dll` used by the client to support Windows authentication.
2. Edit the `be-engine.tra` file and update the `tibco.env.PATH` variable so that it points to the folder where `sqljdbc_auth.dll` resides. For instance:

```
tibco.env.PATH
C:/sqljdbc_4.0/enu/auth/x64%PSP%%BE_HOME%/hotfix/bin%PSP%.
```



JDBC Driver distribution has an x64 or x86 version of that dll. Pick the same version as your BusinessEvents installation.

3. Edit the `studio.ini` file and add the `java.library.path` variable so that it points to the folder where `sqljdbc_auth.dll` resides. For instance:

```
-Djava.library.path=C:/sqljdbc_4.0/enu/auth/x64
```



JDBC Driver distribution has an x64 or x86 version of that dll. Pick the same version as your BusinessEvents installation

4. Use the property `integratedSecurity=true` in the BusinessEvents JDBC resource. For instance:

```
jdbc:sqlserver://sqlserver-group-listener:1433;databaseName=be_
user;integratedSecurity=true
```

5. Replace the script in `<tibco_be_home>/bin/initialize_database_sqlserver.sql` with the following script that creates a user associated to Windows login:

```
use master
go
drop database be_user
go
create database be_user
go
drop login [domain\user]
go
create login [domain\user] from windows with default_database =
be_user
go
use be_user
create user [domain\user] for login [domain\user]
go
```

```
grant control, alter, connect to [domain\user]
go
alter role [db_owner] add member [domain\user]
go
```

SQL Server authentication vs. Windows authentication

As mentioned in the Microsoft documentation, you should use Windows domain logins to access databases that are members of availability groups.

Database users that are associated to domain login added to a database in a primary replica are propagated to secondary databases, and then continue to be associated with the specified domain login.

In the other hand, database users associated with an SQL Server login will be propagated to the secondary databases without a login. The user will not be able to access data from any secondary database.

To overcome this behavior, give the SQL Server login a higher permission at the server level such as `sysadmin` on all the SQL Server replicas, or use Windows login. For Windows login refer to [SQL Server AlwaysOn Availability Groups](#).

Datatype and Driver Information

- The datetime datatype in SQL Server 2005 has the following range: 1/1/1753 to 12/31/9999.
- Microsoft SQL Server 2008 has added a new data type, `datetime2`, which has a date range of 0001/01/01 through 9999/12/31. Therefore, if you are using Microsoft SQL Server 2008, then you can manually change the generated SQL script (DDL) for your backing store, and replace any affected columns' data type from `datetime` to `datetime2`.
- Use the SQL Server JDBC driver, `sqljdbc4.jar`. You can download this driver from: <http://msdn.microsoft.com/en-us/data/aa937724.aspx>

Oracle Database

To install an OCI driver for the Oracle database, follow these steps:

1. Install an Oracle client: preferred as Admin/custom installation.
2. Admin/Custom installation will create the file `tnsnames.ora`.

Configuration while creating the file `tnsnames.ora` should match the database server configuration (database server is running on a remote machine).

3. Install TIBCO BusinessEvents and add the file `ojdbc6.jar` from the client installation to the folder `BE_Home/lib/ext/tpcl`.
4. The connect URL used in the TIBCO BusinessEvents is
`jdbc:oracle:oci:@DBServerHostName:1521:DBInstance`.
 This URL can be found in the file `tnsnames.ora`.
5. If the installation was performed properly, the studio JDBC test connection should show success.

Other Helpful Information

Here are a few helpful points about Oracle database and drivers:

- Use `ojdbc6.jar` drivers file. You can download this file from the following location:
<http://www.oracle.com/technetwork/database/enterprise-edition/jdbc-111060-084321.html>
- Maximum length for an Oracle table name or column name is 30 characters.
- When OCI client is configured with 32-bit client, Oracle might display "No data found" error. To resolve the error, reduce the recovery batch size using the `be.engine.cluster.recovery.batchsize` property.

For example, if you are getting "No data found" error when the batch size is 10000, a batch size of 5000 might resolve the error, that is,

```
<property name="be.engine.cluster.recovery.batchsize"
value="5000"/>
```

Configure the CDD for Special Cases (As Needed)

This section summarizes CDD configuration that may be required to handle special cases as described in [Cases That May Need Additional Setup on page 527](#).

The CDD file is only required during backing store setup if configured for these special cases. Other aspects of CDD configuration are ignored.

Do this aspect of CDD configuration before you use the `be-jdbcdeploy` utility.

Task B Configure the CDD for Special Cases, as Needed

1. In TIBCO BusinessEvents Studio, open the project's CDD file and select Backing Store from the list on the left.
2. In the tree on the left, click **Overrides** and click **Add**. Select the entity or entities you want to customize and click **OK**.
3. In the tree on the left, select the `/uri` entry for each selected entity in turn, and configure the settings on the right:
 - To exclude an entity from the backing store, uncheck the Has Backing Store checkbox.



Memory Only Mode If you configure any entity override Mode setting as Memory Only, then backing store is disabled for that entity.

- To specify a short table name for entities whose table names would exceed the database product's maximum length, check the Has Backing Store checkbox, and enter the name in the Table Name field. See [Ontology Identifiers that Exceed the DBMS Maximum Column Length on page 527](#) for details, and for an alternative way to handle long names.



It is recommended that you specify table names that start with "D_" to match the standard naming convention.

- To specify the length of string properties that exceed 255 characters (that is the actual contents stored in the column is more than 255 characters), check the Has Backing Store checkbox, and in the Properties Metadata section, Max Length setting, specify the expected maximum length for all such properties. See [String Properties that Exceed the DBMS Maximum Column Length on page 528](#) for details.
4. Save the CDD.

Add a JDBC Connection Resource (Now or Later)

Add a JDBC Connection resource to your project and configure it to connect to the backing store. You can do this before or after you configure the database itself. These settings are ignored by the setup utility during the database setup process.

Details below use Oracle 10g database as an example. Adapt the instructions as needed for your database.

See the section [JDBC Connection on page 230](#) for a reference to the fields.



- The value of the CDD Cluster tab > Backing Store > Connection > Max Size field overrides the value of the JDBC Connection Resource Maximum Connections setting.
- For design time features, such as the Test Connection feature, to work correctly, you must add your DBMS product's libraries to the TIBCO BusinessEvents Studio classpath. See the section [Working with External Library and Custom Function Paths on page 8](#).
- For correct runtime operation, see Updating Classpath, Environment Variables, Path, and Copying JAR Files in *TIBCO BusinessEvents Administration*.

Task C Add a JDBC Connection Resource (Now or Later)

1. In TIBCO BusinessEvents Studio, open your project, and open the folder where you keep shared resources and select **New > Other > TIBCO Shared Resources > JDBC Connection**.
2. At the At the New JDBC Connection Wizard dialog, provide a name for the shared connection and click **Finish**. You see the JDBC Connection resource editor.
3. In the JDBC Driver field, select the driver for your database from the selection box on the right, for example, **oracle.jdbc.OracleDriver**. The driver name appears in the box, and a Database URL format appears to the left of the selection box, according to the driver you selected. For the Oracle thin driver, the format is:

```
jdbc:oracle:thin:@<host>:<port#>:<db_instancename>
```

Check the product readme file to ensure you are using the correct database and driver versions.

4. In the Database URL field, configure the provided format. For example:

```
jdbc:oracle:thin:@localhost:1521:ORCL
```

where 1521 is the default port, and ORCL is the default instance name for Oracle Database 10g.

5. In the User Name and Password fields, enter the username and password of the database user (see [Task F, Run the Initialize Database Script as the DBA or System User, on page 543](#)).
6. Save the resource.



Test Connection. After you have configured the database and it is running, you can test the connection. Click the Test Connection button. If the database is running and the details are correct, you see a success message.

Before you can use this feature, you must first add the JDBC driver to the project Properties > Build Path > Java Libraries area in TIBCO BusinessEvents Studio. See note in introduction).

Configure Backing Store Settings in the CDD (Now or Later)

You can do the backing store configuration before or after you configure the database itself. These settings are ignored by the setup utility.

Task D Configure Backing Store Settings and Properties (Now or Later)

Reference Tables

Reference tables for all the properties are documented here:

[CDD Cluster Tab Backing Store Settings Reference on page 457](#)

[CDD Cluster Tab Backing Store Properties Reference on page 461](#)

Configure Database-Related Settings

1. In the CDD editor, select Cluster tab > Object Management > Backing Store.
2. To enable backing store functionality, check the **Enabled** checkbox.
3. In the Database Type field, select **oracle** or **sqlserver**, depending on the type of database you are using.
4. If you select **oracle** then in the Strategy setting, also select the pooling strategy: **oracle** or **jdbc**. Connection pool settings are interpreted differently depending on your choice here, as documented in the reference tables listed above.
5. Choose the database write strategy:
 - To use the cache-aside strategy, check the Cache Aside checkbox.
 - To use the write behind aside strategy, uncheck the Cache Aside checkbox.See Post-RTC and Epilog Handling and Tuning Options in *TIBCO BusinessEvents Architect's Guide* for more on write strategies.
6. If you want to enforce JDBC connection pool settings and properties, check the **Enforce Pools** checkbox. If you do not want to enforce connection pool settings and properties, uncheck the box.

7. Select **Connection** on the left do the following:
 - Select the JDBC Connection shared resource configured for the backing store.
 - Set the minimum, maximum, and initial sizes of the JDBC connection pool as needed.
 - Add additional properties to the Cluster properties sheet as needed (for example, if you are using Oracle Strategy). See [CDD Cluster Tab Backing Store Properties Reference on page 461](#) for details.

Domain Object Overrides

As desired, you can configure CDD settings related to domain objects (entity instances) such as mode, preloading, limited cache options. See [Chapter 28, Domain Objects Configuration, on page 477](#).



If you change settings that enable or disable the backing store for individual entity types after you have created the backing store schema, you must update the schema, as explained in [Updating an Existing Backing Store Schema on page 552](#).

Preventing Database Outages when a Cluster Ceases with Processing

A new feature was added to allow that the Inference engine continues running during a database down period, when all database updates are buffered through the ActiveSpaces caching mechanism. Although the length of the database down period is limited by the available memory for ActiveSpaces datagrid (or by the buffer sizes in case of Coherence).

When database connections are restored, all buffered transactions will be played back and data loss will be avoided.

Required configuration settings:

1. Shared-All persistence with Write-Behind (`Cache Aside=false`)
2. Persistence Mode is set to ASYNC
3. Unlimited cache (both entity and object-table, even if object-table is disabled or `useobjecttable=false`)
4. Preload and Recover all caches (both entity and object-table, even if object-table is disabled or `useobjecttable=false`)
5. `ObjectCacheFullyLoaded` flag is set to 'true' (`<property name="be.engine.cluster.isObjectCacheFullyLoaded" value="true"/>`)

These settings make sure that all the data that Inference engine needs is already cached, and anything that is not found in the cache won't be in the database either. Regardless whether or not database is available, if a 'read/get' does not find the searched entity in cache, database is never queried.

Suggested settings: In addition to the required settings, the following CDD settings are recommended:

1. Cache Agent Quorum = 2 (or more)
2. Number of Backup Copies = 1 (replication of 1 or more)
3. Disable connection checking in Inference agents (prevents exceptions in inference logs)
4. `<property name="be.backingstore.connection.retry.count" value="0"/>` and with cache-aside database connections are not needed except during startup).



If Inference engine is configured to be used as scheduler, then skip this setting.

Cleanup object-table entries marked for deletion during startup (manually or automatically by setting `<property name="be.engine.cluster.cleanup" value="true"/>`)

5. During testing set log-config roles as following (will help debugging)


```
<roles>*:info runtime.service:info kernel.core:debug
backingstore:all jdbcstore:all jdbcstore.impl:all sql.text:all
sql.vars:all</roles>
```

Limitations:

- Scheduler (DB Poller) is database dependent and will not continue executing while database is down.
- If cache server quorum is breached before full recovery, data loss is unavoidable. If every cache servers needs to be shutdown, all the changes done during the database down period will be lost.

Build the EAR File

The EAR file is required when you run the database setup utility. Ontology information in the EAR is used to build tables in the database.

Task E Build the EAR File

Ensure that the project ontology is configured correctly and is complete. Ontology information is required by the database setup utility. Other aspects of TIBCO BusinessEvents Studio configuration are ignored by the database setup utility and can be done before or after you run the utility.

1. In BusinessEvents Studio Explorer, highlight the project name, then from the top menus select **Project > Build Enterprise Archive**.

If you see a message asking you to save all project resources, click Yes. It means an unsaved resource editor is open.

At the Build Enterprise Archive dialog, you can change the EAR file name, specify the location, and complete any other settings as desired.

2. Click **Apply**, then click **OK**. You see messages as the EAR file builds, then you see a message that the EAR file has built correctly:

Initialize the Database and Generate Non-Project Tables

Ensure that you have done all earlier tasks that may pertain to your case (see [Cases That May Need Additional Setup on page 527](#)).

Task F Run the Initialize Database Script as the DBA or System User

This script creates the TIBCO BusinessEvents user and initializes the database.



Running the `initialize_databaseYourDBMS.sql` script deletes the user before creating it again. Running the `create_tables_YourDBMS.sql` drops all database tables before creating them again. This means you can run these scripts again during test phases of your project development, without having to take extra cleanup steps.

The first time you run the `create_tables_YourDBMS.sql` script, you see harmless error or warning messages because there is nothing to delete.

If you are updating the schema for an existing backing store, see [Updating an Existing Backing Store Schema on page 552](#).

1. As desired, change the default TIBCO BusinessEvents user credentials: Open the `initialize_databaseYourDBMS.sql` script for editing and change the default username and password. The documentation uses the default username (`BE_USER`) and password (`BE_USER`)
2. Open a command window in the `BE_HOME/bin` directory (default location of the scripts), and the appropriate command for your DBMS at the prompt:

For Oracle:

```
sqlplus sys_user/sys_user_password@SID
@initialize_database_oracle.sql
```

Type `exit` to exit and commit.

If you are using SQL Server, use this instead:

```
osql -S Your-Server -U sys_user -P sys_user_password -n -i
initialize_database_sqlserver.sql
```

This script creates the TIBCO BusinessEvents database user. This user must be used to run the other scripts. You see messages like the following:

```
User dropped.
User created.
Grant succeeded.
```



Using your database product, you can configure additional users to access the database, in addition to this user.

Task G Run the Create Tables Scripts as the TIBCO BusinessEvents User

Log on as the TIBCO BusinessEvents user, `BE_USER` by default and run a script to create non-project specific tables.

1. Open a command window in the `BE_HOME/bin` directory (default location of the scripts), and type the appropriate command for your DBMS at the prompt:

For Oracle:

```
sqlplus BE_USER/BE_USER@SID @create_tables_oracle.sql
```

Type `exit` to exit and commit.

If you are using SQL Server, use this instead:

```
osql -S Your-Server -d Your-DB -U BE_USER -P BE_USER -n -i
@create_tables_sqlserver.sql
```

Use the credentials defined in the `initialize_database_oracle.sql` or `initialize_database_sqlserver.sql` files. By default those are: username `BE_USER`, with password `BE_USER`.

You see various harmless error messages the first time you run the script.



TIBCO BusinessEvents Studio provides a wizard to configure the required properties and generate the project-schema-specific SQL scripts. The wizard is an alternative approach to perform [Task H](#). See [Generating Scripts Using the JDBC Deployment Wizard on page 545](#) for details.

Generate the Project-Schema-Specific SQL Scripts

You can do this task using a TIBCO BusinessEvents Studio wizard, or you can do it manually.



Before you Begin

- Ensure that you copied your JDBC drivers file to `BE_HOME/lib/ext/tpcl` (or other location in your class path). See [Task A, Ensure DBMS Software and Drivers are in Place, on page 531](#).
- Ensure the DBMS is started.

Generating Scripts Using the JDBC Deployment Wizard

You can generate the project-schema-specific SQL scripts using the **JDBC Deployment** wizard in TIBCO BusinessEvents Studio.

Task H Generate the Project-Schema-Specific SQL Scripts (with Wizard)

1. In TIBCO BusinessEvents Studio Explorer, right-click the name of the project for which are creating SQL scripts and select **Export > TIBCO BusinessEvents > JDBC Deployment**. Click **Next**.
2. You see the **Generate JDBC deployment scripts** wizard. Complete the values as follows.

Table Creation
Scripts

Field Name	Description
Database Type	Select the type of database you are using from the drop-down list. Choose one of: oracle and sqlserver The default value is oracle
Generate ANSI Scripts	Check this checkbox to use ANSI compatible SQL types during script generation. Default setting is checked.
Cluster Deployment Descriptor	Browse to and select the CDD file you want to use for the project.

Field Name	Description (Cont'd)
Output Directory	Browse to and select a directory where the scripts are to be generated, for example, <i>BE_HOME/bin</i> (This directory is used if you generate files manually.)
Output Script Name Prefix	Enter a prefix for the output script filenames. For example, if you enter <i>acme</i> , the following scripts are generated: <i>acme.sql</i> , <i>acme.aliases</i> , <i>acme_alter.sql</i> , <i>acme_cleanup.sql</i> , <i>acme_delete</i> , and <i>acme_remove.sql</i> .

3. Do one of the following:

- If you are creating a new JDBC backing store, click **Finish**. The scripts are generated in the location you specified. Information about script generation is also printed to the Console tab.
- If you are migrating an existing JDBC backing store, click **Next**. You see the Generate Migration SQL Scripts page. See [step 4](#) for details.

Migration Scripts 4. At the Generate Migration SQL Scripts page, complete the values as follows.

Field Name	Description
Generate Migration Scripts	Check this checkbox to generate SQL scripts to migrate an existing backing store. Once checked, the rest of the fields on this page are enabled. The default value is unchecked.
Connection Configuration	Browse to and select the JDBC Connection shared resource used to connect to the existing backing store.
Database URL	Enter the database URL that points to the existing backing store.
Database Username	Enter the database username that was used when setting up the existing backing store.
Database Password	Enter the database password that was used when setting up the existing backing store.
Database Schema Owner	Enter the database schema owner for the existing backing store.

Field Name	Description (Cont'd)
Test Connection	Click Test Connection to test if a success

Generating Scripts Manually

This section explains how to generate scripts using a manual procedure, instead of using the TIBCO BusinessEvents Studio wizard.



If you are using SQL Server: Before you execute `be-jdbcdeploy`, open the `BE_HOME/bin/be-jdbcdeploy.tra` file for editing. Specify `sqlserver` in the last line:

```
java.property.jdbcdeploy.database.type [oracle | sqlserver]
```

This step requires the EAR file for the project. The database utility uses the project ontology information from the EAR file.

1. Open a command window and navigate to `BE_HOME/bin`.
2. Run `be-jdbcdeploy.exe` using a command with the following format:

```
be-jdbcdeploy [-h] [-p property file] [-o schema output file] [-c CDDpath] EAR Path
```

For example:

```
be-jdbcdeploy -o acme -c D:/myproj/acme.cdd D:/ears/acme.ear
```

The generated scripts appear in the directory where you run the executable. For example, if you provide the schema output filename `acme`, you would see files called `acme.sql`, `acme.aliases`, `acme_alter.sql`, `acme_cleanup.sql`, `acme_delete`, and `acme_remove.sql`.

The user-defined part of the database schema is in the schema output file (`yourname.sql`) as schema definition commands. In [Task J](#) you run this script (together with provided scripts) to build the schema in the database.

The options are explained in the following table:

Option	Description
<code>-p, /p,</code> <code>-property,</code> or <code>/property</code>	Optional. Specifies the property file. If not specified, the default property file is used, that is, <code>be-jdbcdeploy.tra</code> in the current directory.

Option	Description
-o	<p>Required. Specifies the schema output filename for deployment.</p> <p>Tip If you specify a directory path, the backing store scripts are generated in the specified directory and the last element of the path is taken as the schema output filename.</p>
-c	<p>Optional. Specifies the absolute path to the CDD file. See Cases That May Need Additional Setup on page 527 to understand when the CDD file is needed.</p>
<i>EARpath</i>	<p>Required. The last option is always the EAR file path.</p>
-h, /h, or /help	<p>Displays this help.</p>
-ansi, /ansi	<p>Optional. If set to true, ANSI compatible SQL types are used during script generation.</p> <p>Allowable values are true and false. For ANSI compatible databases, it is set to true by default.</p>

Check the Aliases File, Run the Project Schema Script

For every entity, property, or state machine whose database identifier name exceeds the database maximum length, a table name entry is created in the generated *yourname*.aliases file (for example, acme.aliases). See [Ontology Identifiers that Exceed the DBMS Maximum Column Length on page 527](#) for more information, and for an alternative way to specify short table names.

It's a good idea to check the aliases file for entries, even if the TIBCO BusinessEvents names are not very long. The length of the generated database table names is not easy to predict.

Optionally, you can edit the file to provide more meaningful names.



It is recommended that you keep the aliases file for future reference. If the project ontology changes after the backing store has data in it, you must also update the database schema to match the new schema (as explained in [Updating an Existing Backing Store Schema on page 552](#)). If you modified the generated aliases, you must use the same aliases again when you update the schema, to preserve those columns and their data.

Key word
mapping file

Entries in the key word mapping file are also added to the aliases file so you can replace the key word aliases with project-specific ones, as desired (generally in a second pass). For details see [Task K, If Needed — Map Key \(Reserved\) Words to Aliases, on page 550](#).

Task I Check the Aliases File and Modify Aliases as Desired

1. Open the *yourname*.aliases file for editing.
2. Replace any aliases as desired with more meaningful short names. Make sure that each name is unique. It's a good idea to leave any system generated prefixes or suffixes in place for consistency of names across the database.
3. Re-run the `be-jdbcdeploy` tool, using the same parameters as before. (For details see [Task H, Generate the Project-Schema-Specific SQL Scripts \(with Wizard\), on page 545](#).) This time, the aliases you created are used.

Task J Run the Project Schema Script (as BE_USER)

In this step, you log on as the user you created and run a script to create the project related part of the database schema.

1. Open a command window in `BE_HOME/bin` and run the *yourname*.sql script. (for example, `@acme.sql`). Use the user `BE_USER`, password `BE_USER` (or whatever username and password you set in the script in [Task F](#)).

For Oracle:

```
sqlplus BE_USER/BE_USER @yourname.sql
```

For SQL Server:

```
osql -S Your-Server -d Your-DB -U BE_USER -P BE_USER -n -i @yourname.sql
```

If there are no errors, your database tables are now configured for use. If there are errors you may need to add some mappings to the key word mapping file.

Task K If Needed — Map Key (Reserved) Words to Aliases

Complete this task only if you saw errors after completing [Task J, Run the Project Schema Script \(as BE_USER\)](#). Such errors are caused when your project ontology uses terms that are key words (reserved terms) in the DBMS you are using. You must map these terms to an alias in the keyword mapping file.

1. Edit the `BE_HOME/dbkeywordmap.xml` file to add entries. Below is the format followed by an example:

```
<keyword name="dbKeyword" mapname="nonDbKeyword" />
<keyword name="start" mapname="start_" />
```

2. Repeat [Task H, Generate the Project-Schema-Specific SQL Scripts \(with Wizard\)](#), on page 545, and tasks following as needed.



Providing Project-Specific Key Word Aliases When you repeat [Task H](#), the new key words are added to the `yourname.aliases` file. You can create project-specific aliases for the key word mappings as desired. Then repeat [Task H](#) again and continue.

Note that you must generate the SQL scripts a minimum of three times if you add keyword mappings to the aliases file — it might be more because you may not catch all errors at once. For example, if there are multiple keyword clashes in one table, only the first are reported. Perform this loop until no more errors occur. A summary of the basic routine is as follows:

1. *Generate the SQL scripts and run them (as explained in the procedures).*
2. *Errors occur due to key word clashes, so you add the appropriate key word mapping entries to the key word mapping file.*
3. *Generate the SQL scripts again.*
4. *To use project-specific aliases for the keyword mappings (Optional):*
 - a. *Edit the aliases file entries for the key word mappings.*
 - b. *Generate the SQL scripts again.*
5. *Run the SQL scripts to create the backing store.*

Next Step (As Needed)

Complete the TIBCO BusinessEvents Studio project configuration tasks if you have not already done so. These tasks can be done before or after database setup. See the following sections:

- [Add a JDBC Connection Resource \(Now or Later\) on page 537](#)
- [Configure Backing Store Settings in the CDD \(Now or Later\) on page 539](#)



Update your schema if your ontology changes, or if you want to include or exclude different entities in the backing store. See [Updating an Existing Backing Store Schema on page 552](#) for details.

Updating an Existing Backing Store Schema

If you change the project ontology, that is, if you create, alter or delete a concept or an event, you must update the backing store schema so it matches the updated ontology. In the case of changes in project ontology, you must update the backing store schema before you deploy the updated project.

You may also wish to change which entities are excluded from the backing store using CDD settings (see [Excluding Entities from the Backing Store on page 528](#) for more information. This change does not require project redeployment. It requires that the updated CDD file is copied to all runtime machines.

Examine the alter script before you run it. The section [What the Schema Update Utility Can and Cannot Handle Automatically on page 553](#) provides more information.

Preparing to Run and Running the Schema Update Utility

1. To prepare for the update, do the following:
 - Gracefully shut down the deployed application (all agents including cache agents).
 - Back up your existing database.
 - Generate the updated EAR file for the modified project.
 - If you modified aliases when you created the schema, locate the *yourname.aliases* file you used. It will help you to modify those aliases in the newly generated file, so they match.
2. Open the `be-jdbcdeploy.tra` file for editing and set the following properties:

```
be.jdbc.schemamigration.url=DbURL  
be.jdbc.schemamigration.user=username  
be.jdbc.schemamigration.pswd=password
```

- Use the database URL that points to the existing backing store. See [Add a JDBC Connection Resource \(Now or Later\) on page 537](#) for example URLs.
- Use the same username and password you used when setting up the backing store. See [Task F, Run the Initialize Database Script as the DBA or System User, on page 543](#).

These properties enable the program to compare the schema of the existing database with the ontology in the project EAR file, and generate the alter script.

3. Log on as the user name you specified in [Task F, Run the Initialize Database Script as the DBA or System User, on page 543](#). For example,
4. Run the `be-jdbcdeploy` utility as explained in [Task H, Generate the Project-Schema-Specific SQL Scripts \(with Wizard\), on page 545](#), using the *updated* EAR file.
5. If any of the new or changed definitions result in entries in the `yourname.aliases` file, and you want to change the provided aliases, follow instructions in [Task I, Check the Aliases File and Modify Aliases as Desired, on page 549](#). If you modify aliases, remember to generate the scripts again so the modified aliases are used.



You must use the same aliases that you used before. If any were modified when the schema was created, you must modify them the same way when updating the schema. It can be useful to refer to the original aliases file.

6. Examine the generated `yourname_alter.sql` script and modify as needed so you only run statements for changes you want to make. See [What the Schema Update Utility Can and Cannot Handle Automatically on page 553](#) for details.
7. Run the `yourname_alter.sql` script.

Your database tables are now configured for use.

What the Schema Update Utility Can and Cannot Handle Automatically

You must examine the alter script before you run it. Decide what changes to make manually and what changes to make using the script, taking into account the kind of data in the tables. Entries that could result in data loss are commented. Remove or comment entries for changes you will make manually.

Adds

The schema migration utility handles addition of entity types and attributes. New entity types and attributes are added to the database schema.

Changes (Drop and Add) — Assess individually

The utility handles changes to attributes (entity properties) as `DROP` and `ADD` operations. However, `DROP` operations are commented in the script to avoid data loss.

If a column is empty, or you do not want to keep the data they contain, you can enable the `DROP` operation and let the utility handle the change.

If the column contains data that you want to keep, then make the change manually using an appropriate database tool. For example, you can change the data type of a column from string to double without loss of data, as long as all the column values are numeric values.

Entity Deletions

If an entity is deleted from the TIBCO BusinessEvents Studio project, the corresponding tables are not dropped from the database schema. Existing data is not lost. Deleted entities are not mentioned in the alter script. Manually keep track of and delete such tables as needed.

Attribute Deletions

The schema update utility does handle deletion of entity attributes. SQL statements for deleted attributes are generated but they are commented. Examine the alter script and enable these commands if you want to execute them. Note that existing data is lost when you drop an attribute.

Example Alter Script

Below is an example *yourname_alter.sql* script.

Property type change	<pre>-- ##### WARNING : Non-alterable Ontology changes found. Please see following errors. Manual schema-migration is required. --* For Concept Concept1 field PROPERTY_1 type changed from VARCHAR2 to LONG -- ALTER TABLE D_Concept1 DROP (Property_1); ALTER TABLE D_Concept1 ADD (Property_1 numeric(19));</pre>
New table	<pre>DROP TABLE D_Book_rrf; CREATE TABLE D_Book_rrf (pid numeric(19), propName char varying(255), id\$ numeric(19) not null);</pre>
New property	<pre>-- ALTER TABLE D_MyConcept DROP (FOLDER_1); ALTER TABLE D_MyConcept ADD (Folder_0 char varying(255));</pre>

Handling Entities Deleted from a Write-Behind Backing Store



TIBCO BusinessEvents Express This section relates to Cache OM functionality and does not apply to TIBCO BusinessEvents Express edition.

This section pertains only if write-behind strategy is used. See Post-RTC and Epilog Handling and Tuning Options in *TIBCO BusinessEvents Architect's Guide* for more information about database write strategies, and see [CDD Cluster Tab Backing Store Settings Reference on page 457](#) for details on the settings.

When write-behind strategy is used, when entities are deleted from the cache, they are marked for deletion at system startup. An application startup process removes the deleted entities from the backing store.

Depending on the number of deleted entities, the cleanup step can make system startup unacceptably slow.

Manual Deletion of Deleted Entities

To deal with slow startup, you can run a script manually, and eliminate the need to run the cleanup action at system startup. The script is called `yourname_delete.sql` and it is generated when you create the backing store with the `be-jdbcdeploy` utility.

Disabling Deleted Entity Cleanup at System Startup

You can also disable the cleanup action performed at system startup and perform all cleanup manually, using the deletion script. To do so, add the following property to the CDD file and set it to false:

```
be.engine.cluster.cleanup
```

Backing Store Table Reference

The backing store uses relational tables and SQL data types for ease of maintenance. The SQL (DDL) scripts use ANSI SQL type definitions (where supported by the target DBMS product).

Each ontology type in the backing store has its own primary table and zero or more second-level tables. There are only two levels of tables, which makes the database easier to manage and easier to understand. Because the backing store adheres to SQL standards and a straight-forward structure, standard database tools can be used to view backing store data.

Primary Tables

Primary tables contain only primitive properties such as the following:

Property	Note
cacheId	Entity version number (starts with 1)
time_created\$	Time when the entity was created
time_last_modified\$	Time when the entity was last modified
parent\$_id\$	Id of the parent for contained concepts
id\$	Unique Id of the entity (must be unique across all entities)
extId\$	Unique (or null) extId assigned
states	Always set to 'C' meaning 'Created' (reserved for future use)

Secondary Tables

Secondary tables are used for complex properties, that is, arrays, properties with history, and concept relationship properties. Each array and history-enabled property has a separate table. Only primitive properties are stored in the primary table.

Table 63 Secondary table structure

property Type	Column	Description
Array	pid\$	Parent ID
	valPid\$	Array index
	val	Item's value
History	pid\$	Parent ID
	howMany	Number of history items
	timeIdx	Item's time stamp
	val	Item's value
Array with History	pid&	Parent ID
	valPid\$	Array index
	howMany	Number of history items
	timeIdx	Item's time stamp
	val	Item's value

Reverse Reference Tables

Each concept also has a reverse reference table. This table's name contains the concept name and ends with the characters `_rrf$`. It has these columns:

Column	Description
pid\$	Parent ID from the main concept table
propertyName\$	Property name (field) from the referencing concept.
id\$	Identifier (id\$) of the referencing concept.

Class-to-Table Mapping

This table contains the mapping between class names and table names, and the mapping between complex property field names and secondary table names.

for example:

```
'be.gen.Ontology.DeleteVerifyEvent', 'D_DeleteVerifyEvent'  
'be.gen.Ontology.Treatment', 'D_Treatment'  
'be.gen.Ontology.Treatment', 'rrf$', 'D_Treatment_rrf$'  
'be.gen.Ontology.BaseAlert', 'treatments',  
'D_BaseAlert_treatments'
```

Chapter 31 **Introduction to WebStudio**

This chapter introduces the TIBCO BusinessEvents WebStudio application and its server component, Rules Management Server (RMS).

Topics

- [Introduction to TIBCO BusinessEvents WebStudio, page 560](#)
- [Introduction to Rules Management Server \(RMS\), page 564](#)
- [WebStudio and RMS User Workflow, page 566](#)

Introduction to TIBCO BusinessEvents WebStudio

TIBCO BusinessEvents WebStudio is a TIBCO BusinessEvents online component consisting of a client, WebStudio, and a server, Rules Management Server (RMS).

WebStudio Component

TIBCO BusinessEvents WebStudio is an online component which allows business users to create/manage business rules in a web browser. In TIBCO BusinessEvents WebStudio user defines an executable rule (business rule) based on the rule template and rule template view defined by developer in TIBCO BusinessEvents Studio. Similar to TIBCO BusinessEvents Decision Manager, a decision table in TIBCO BusinessEvents WebStudio is defined using the virtual rule function. Rule templates, rule template views, and virtual function are created in TIBCO BusinessEvents Studio by developers and stored in RMS repositories.



The decision tables are available only if TIBCO BusinessEvents Decision Manager add-on is installed.

Rules Management Server (RMS) Component

Rules Management Server (RMS) is a server-based component that manages the lifecycle of WebStudio artifacts and WebStudio user access.

Using a menu of options in TIBCO BusinessEvents WebStudio, RMS allows business users to check out projects and check them in, subject to an approval process that RMS also provides. The workflow process can be customized. For a more detailed overview see [WebStudio and RMS User Workflow on page 566](#).

Before RMS can be used, some configuration is required, as explained in [Introduction to Rules Management Server \(RMS\) on page 564](#).

Technical User Tasks

As a technical user you configure RMS and the TIBCO BusinessEvents WebStudio components, and set up project resources required for business users to create decision tables or business rules. You also deploy decision tables or business rules when they are ready for use, either in an enterprise archive (EAR) file, or as class files. Class files can be deployed for use at startup, or hot deployed. They can also be unloaded from a running engine. You also define access control settings and can act as an approver to ensure that decision tables and business rules created by business users are appropriate for use.

Virtual Rule Functions and Decision Tables

At designtime, technical users add *virtual rule functions* (VRFs) to a TIBCO BusinessEvents project. A VRF has no body, similar to a Java interface. Its implementation is provided using decision tables authored in TIBCO BusinessEvents WebStudio. VRFs are used in the TIBCO BusinessEvents project like any other rule function; they can be called from rules or other rule functions. Here is a simple example:

```
/**
 * @description
 */
virtual void rulefunction Virtual_RF.Applicant_VirtualRuleFunction {
    attribute {
        validity = ACTION;
    }
    scope {
        Concepts.Applicantapplicant;
        Events.ApplicationReceivedapplicationreceived;
    }
    body {
    }
}
}
```

In TIBCO BusinessEvents WebStudio, business users add decision table resources to VRFs. The decision table provides the body to the VRF, also known as the VRF implementation.

One VRF can have more than one decision table. If a VRF has more than one decision table, functions in TIBCO BusinessEvents determine how the tables are used.

Rule Template and Business Rule

At designtime users add rule template to a TIBCO BusinessEvents project. Rule template is a specialized rule that contains a pre-conditions section which defines the pre-conditions which must be met in addition to the conditions defined by the business user in a business rule.

The Action Context section of a rule template defines all possible actions that can be taken by a business rule (after all conditions are met). Only the action context statements that the WebStudio user selects and defines as commands in the business rules are actually taken (depending on rule evaluation at runtime). For a business rule execution to succeed, a business rule definition must include all Actions and the order should be maintained. Completing the definition of an

action is a WebStudio user task. If bindings are used (and a view) then in WebStudio, the business rule writer only has to enter the binding values to complete the definition. Action context statements are of three types: create, modify, and call, plus arbitrary actions.

Business User Task

Rule Building with Business Rules

As a business user, you check out projects from RMS, build business rules, and submit them for approval.

Business rules provide builder or a user-friendly HTML form to build rules. In builder you specify the conditions in the “when” section and the actions in the “then” section. You can use artifacts and supported operators for building conditions and actions. In user-friendly HTML form, you fill up the values for the conditions and actions.

Rule Building with Decision Tables

As a business user, you check out projects from RMS, build decision tables, and submit them for approval.



Decision tables are available in TIBCO BusinessEvents WebStudio in edit mode only if you have TIBCO BusinessEvents Decision Manager add-on installed on your system.

Decision tables provide a graphical way to build complex business rules. You create table columns by dragging and dropping predefined properties onto the decision table framework. The properties belong to ontology resources defined in the TIBCO BusinessEvents project. However, you can only use the properties specified in the VRF. Columns can be created in other ways too. You then define threshold values (conditions) and actions in the cells of the table. Each row can be thought of as one rule in a table made up of many rules. The individual rules are often straightforward, as in the following examples.

Three Rule Conditions	<pre>Person.age < Max(20, Parent.age) Person.creditscore >= Math.function(...) Person.gender == "female"</pre>
--------------------------	--

Three Rule Actions	<pre>Application.status = "ACCEPTED" Application.credit = 4000 sendNotification()</pre>
-----------------------	---

However, one decision table can consist of hundreds, even thousands of rules each of which is executed only when its specific conditions are satisfied.

Exception Tables

Each decision table can optionally have another table known as an exception table. The purpose of the exception table is purely organizational: it enables you to separate the business logic of the main decision table (added by business users) from any non-business logic (generally added by technical users). For example, in the exception table, you could capture situations where fields are blank or contain invalid values, and define actions that send notifications or set return values. The rows of a decision table plus the rows of its exception table are considered in an RTC. If you prefer, you can put non-business logic in the main table instead of using an exception table.

Table Analyzer

The Table Analyzer feature analyzes decision tables and reports problems, such as uncovered ranges for conditions, uncovered domain entries, different set of actions for identical conditions. Table Analyzer creates a sparse matrix data structure representing an optimized form of decision tables in memory.

Introduction to Rules Management Server (RMS)

RMS is a lightweight rules management server component for managing the repository of TIBCO BusinessEvents projects. RMS provides an easy, secure, and scripted deployment lifecycle. RMS is supported on all platforms that support TIBCO BusinessEvents. Users can access RMS remotely. RMS does not have to be installed on users' machines.

Under certain circumstances, use of RMS is optional. (see *TIBCO BusinessEvents Decision Manager User's Guide*).

Before you can use RMS some basic configuration may be required. You also set up an RMS project for each TIBCO BusinessEvents project that includes decision tables or business rules. See [Chapter 32, RMS and RMS Projects Configuration, on page 567](#).

RMS enables business users to check out projects. The necessary project resources loaded into the cache. If you have the TIBCO BusinessEvents Decision Manager add-on installed, then the artifacts are also copied into the clients machine. When users are finished working with the projects, they submit the business rules and/or decision tables to RMS for approval. (See [WebStudio and RMS User Workflow on page 566](#) for a more detailed overview.)

RMS enables technical users to do the following:

- Configure authentication and access control as needed. See *TIBCO BusinessEvents Administration* for authentication and access control configuration.
- Approve or reject commit requests from Decision Manager and TIBCO BusinessEvents WebStudio users, check on the status of all such requests, keep track of all project versions, and view audit trails of status changes during workflow execution.
 - [Chapter 36, Working with the Approval Workflow in RMS, on page 625](#).
- Generate project Enterprise Archive (EAR) files or individual decision table classes or rule template instance file for business rule, for deployment, and deploy them.
 - [Chapter 37, Deployable Files Generation and Deployment for WebStudio, on page 631](#).

RMS is implemented using TIBCO BusinessEvents. Knowledgeable TIBCO BusinessEvents users can customize it. The RMS projects are located in the *BE_HOME/examples/standard/WebStudio* directory, you can configure the RMS to take projects from SVN location as well. Note that documentation is provided for the product as shipped. If you customize RMS some documentation may not apply.

WebStudio and RMS User Workflow

This section describes how rules authored in TIBCO BusinessEvents WebStudio are made available to a TIBCO BusinessEvents application.

Step 1: TIBCO BusinessEvents Designtime

A TIBCO BusinessEvents user creates a TIBCO BusinessEvents project, adding the ontology, and writing rules that make use of virtual rule functions (VRFs) and Rule Templates. At designtime, VRFs have a signature but no body.

RMS Project Setup

The TIBCO BusinessEvents project for RMS is stored at the location defined by a property in the RMS server configuration. The RMS project requires an access control file.

The RMS server must be running so that the login and workflow actions are available in the TIBCO BusinessEvents WebStudio component.

Step 2: TIBCO BusinessEvents WebStudio

A business user starts the TIBCO BusinessEvents WebStudio component, logs on to RMS, and checks out the project. The business user creates one or more decision tables or business rules and saves the modified project locally, then commits them for approval.

Step 3: Approval

An RMS user working in TIBCO BusinessEvents WebStudio receives the request and reviews the checked-in artifacts and then approves or rejects them. The approved artifacts are available for subsequent checkouts or updates.

Step 4: Generating Deployable Files

An RMS user generates deployable files for resources that are ready for deployment. You can generate EAR files or class files.

Step 5: Deployment

EAR files are deployed in the usual way, as explained in *TIBCO BusinessEvents Administration*. Class files for decision tables and rule template instance files for business rules are deployed by placing them in a configured location recognized by the TIBCO BusinessEvents engines at startup. They can also be hot-deployed.

Chapter 32 **RMS and RMS Projects Configuration**

This chapter explains how to add a project to the Rules Management Server (RMS). The chapter also explains the configuration tasks for RMS, for RMS projects, and for TIBCO BusinessEvents WebStudio.

Topics

- [Adding a Project to RMS, page 568](#)
- [Configuring RMS Server Properties, page 569](#)
- [Enabling Remote Connection to RMS from WebStudio, page 571](#)
- [Enabling SSL Authentication in RMS, page 572](#)
- [Adding Shared JDBC Connection in Runtime, page 574](#)
- [RMS Server Configuration Property Reference, page 575](#)

Adding a Project to RMS

TIBCO BusinessEvents projects, with VRF and rule templates, needs to be added to the RMS repository to make them accessible in the TIBCO BusinessEvents WebStudio. All RMS projects exist in a repository under a *root location* directory and their access control list (ACL) file exist under a *RMS config* location, which are `BE_HOME/examples/standard/WebStudio` and `BE_HOME/rms/config/security` respectively, as shipped. The root location and RMS config location are configurable using `ws.scs.rootURL` and `ws.projects.acl.location` properties in the `RMS.cdd` file, respectively.

Prerequisite

A TIBCO BusinessEvents project created in TIBCO BusinessEvents Studio.

Procedure

1. Under the RMS project root location directory, create a directory with a name appropriate for the project.

The root location is configurable. See [Configuring RMS Server Properties on page 569](#).

2. Copy the TIBCO BusinessEvents Studio project contents that will be used in Decision Manager add-on or in TIBCO BusinessEvents WebStudio to the project directory.
3. If the project contains custom functions, create the *ProjectName* directory under the directory identified by the `ws.projects.customLib.location` property in the `RMS.cdd` file and then place the custom function jars under it. See [RMS Server Configuration Property Reference on page 575](#) for more information on the property.
4. Copy the ACL file for the project (`RMSProjectName.ac`) to the RMS config directory. See *TIBCO BusinessEvents Administration* for authentication topics.

The RMS config location is configurable. See [Configuring RMS Server Properties on page 569](#).

Configuring RMS Server Properties

RMS works out of the box on a local machine. Configuration is required to enable access by remote TIBCO BusinessEvents WebStudio and the Decision Manager clients.

You may need to make other changes to the server configuration. For example, if you change the server location or the location of the RMS project repository, you must update the RMS server properties accordingly.

To configure the RMS server properties, edit the `RMS.cdd` file. You cannot edit a CDD file in TIBCO BusinessEvents Studio outside of its project context. To edit the CDD in TIBCO BusinessEvents Studio, import the BRMS project into your workspace. After that copy the `RMS.cdd` file to its installed locations, in the following way.



Whenever you change the `RMS.cdd` file you must restart the RMS server so that it uses the updated values.

See [RMS Server Configuration Property Reference on page 575](#) for information about each property.



- As with any procedure that changes files, first backup any files that could be affected. The `RMS.cdd` file is located in two places, in the product as shipped:
 - `BE_HOME/rms/project/BRMS/Deployment/RMS.cdd`
 - `BE_HOME/rms/bin/RMS.cdd`
- After you update the `RMS.cdd` file in your workspace, overwrite the existing files with the updated ones, in the above locations (or the ones currently in use if different).
- Ensure that the CDD files in the BRMS project and the RMS startup directory are kept in sync.

Procedure

1. In TIBCO BusinessEvents Studio, choose **File > Import > Existing Projects into Workspace**.
2. Ensure the Copy the projects into workspace check box is selected.
3. Select the following project:
 - `BE_HOME/rms/project/BRMS`
4. In Studio Explorer, double-click **RMS.cdd** to open it in the CDD editor.

5. In the CDD editor Cluster tab, click **Properties**. In the Configuration panel on the right, expand groups to see the individual properties. Complete the values as explained in [RMS Server Configuration Property Reference on page 575](#).
 - [Enabling Remote Connection to RMS from WebStudio on page 571](#)
 - [Enabling SSL Authentication in RMS on page 572](#)
 - [Adding Shared JDBC Connection in Runtime on page 574](#)
 - [RMS Server Configuration Property Reference on page 575](#)
6. When you have finished editing, save the `RMS.cdd` file.
7. In the file system, copy the `RMS.cdd` file from your workspace to the BRMS project and to the RMS server startup location:
`BE_HOME/rms/project/BRMS/Deployment/RMS.cdd`
`BE_HOME/rms/bin/RMS.cdd`
8. Restart the RMS server.

Enabling Remote Connection to RMS from WebStudio

As shipped, the `tibco.clientVar.Webstudio/hostname` property is set to `localhost`. This setting enables the product to be used on a single machine. For production settings specify the RMS server hostname and port to ensure that TIBCO BusinessEvents WebStudio users can connect to RMS from remote machines.

Procedure

1. Import the BRMS project into your workspace and open the `RMS.cdd` file for editing. See [Configuring RMS Server Properties on page 569](#) for details.
2. In the CDD editor Processing Units tab, click **WS-Inference**. In the Properties panel on the right, expand the WS property group.
3. Specify the TIBCO BusinessEvents WebStudio hostname and port using the following properties:

```
tibco.clientVar.Webstudio/hostname  
tibco.clientVar.Webstudio/port
```

4. Save the `RMS.cdd` file
5. In the file system, copy the `RMS.cdd` file from your workspace to the BRMS project and to the RMS server startup location:

```
BE_HOME/rms/project/BRMS/Deployment/RMS.cdd  
BE_HOME/rms/bin/RMS.cdd
```

6. Restart the RMS server.

Enabling SSL Authentication in RMS

The SSL authentication process uses certificates that are issued by a certificate authority. The same process applies if the certificates are issued by a certificate generation utility or if self-signed certificates are used.

To establish a SSL connection an application acting as an SSL client contacts an application acting as an SSL server. There are two ways in which SSL authentication can be made:

- One-way SSL authentication
- Two-way SSL authentication

Enabling One-way SSL Authentication in RMS

One-way SSL authentication enables the application operating as the SSL client, which in this case is web browser, to verify the identity of the application operating as the SSL server (RMS). The SSL-client application is not verified by the SSL-server application.

Procedure

1. Import the BRMS project in the TIBCO BusinessEvents Studio and copy it into your workspace. Now open the `RMS.cdd` file for editing. See [Configuring RMS Server Properties on page 569](#) for details.
2. In the CDD editor Processing Units tab, click **default**. In the Properties panel on the right, you can see the `be.channel.deactivate` property.
3. Remove the `"/WebStudio/Core/Channels/WS_CH_Secure_WebstudioChannel"` value from the `be.channel.deactivate` property.
4. Save the `RMS.cdd` file
5. In the file system, copy the `RMS.cdd` file from your workspace to the BRMS project and to the RMS server startup location:
`BE_HOME/rms/project/BRMS/Deployment/RMS.cdd`
`BE_HOME/rms/bin/RMS.cdd`
6. Restart the RMS server.

To verify, type the URL `"https://localhost/WebStudio/"` in a web browser and press **Enter**. No port is required as default SSL port is used. The web browser notifies you about the untrusted certificate from the SSL server.

Enabling Two-way SSL Authentication in RMS

In two-way SSL authentication, the SSL client application, which in this case is a web browser, verifies the identity of the SSL server application (RMS), and then the SSL server application verifies the identity of the SSL-client application.

Procedure

1. Follow the steps mentioned in the [Enabling One-way SSL Authentication in RMS on page 572](#) to enable the SSL server verification, but do not restart RMS server yet.
2. In Studio Explorer, double click **BRMS > WebStudio > Core > Transports > WS_TR_Secure_WebstudioHTTPConnection.sharedhttp** to open it in the resource editor.
3. In the resource editor, click **Configure SSL**. The SSL Configuration for HTTPS Connections window is displayed.
4. In the SSL Configuration for HTTPS Connections window, select the **Requires Client Authentication** checkbox and click **OK**.
5. In the BusinessEvents Studio, regenerate the BRMS project EAR file (RMS.ear) and save it at `[BE_HOME]/rms/project/`. See [Building an Enterprise Archive \(EAR File\) on page 22](#) for more details.
6. Restart the RMS server.
7. Open your web browser and import the `WSClientStore.p12` certificate file in the web browser from the location `[BE_HOME]/rms/config/security/`. Refer to your web browser's help on how to import certificate in it.

To verify, type the URL "http://localhost/WebStudio/" in the same web browser and press **Enter**. No port is required as default SSL port is used. The web browser displays a confirmation window to select your SSL client certificate and send it to the SSL server. After that the web browser notifies you about the untrusted certificate from the SSL server.

Adding Shared JDBC Connection in Runtime

To setup the shared JDBC connection to the backing store for RMS at runtime, you do not require to generate new RMS.ear file as identified in [Add a JDBC Connection Resource \(Now or Later\) on page 537](#). Instead, once your backing store is setup, you can edit the RMS.cdd file to add the JDBC connection to the backing store from RMS.

Procedure

1. Import the BRMS project in the TIBCO BusinessEvents Studio and copy it into your workspace. Now open the RMS.cdd file for editing. See [Configuring RMS Server Properties on page 569](#) for details.
2. In the CDD editor Cluster tab, click the **Properties** in the Cluster pane and expand **node** under the Configuration pane.
3. Expand the **BackingStore** property group and edit the value for backing store properties as required. See [Backing Store Property Group on page 582](#) for more details on the properties.
4. Save the RMS.cdd file
5. In the file system, copy the RMS.cdd file from your workspace to the BRMS project and to the RMS server startup location:
`BE_HOME/rms/project/BRMS/Deployment/RMS.cdd`
`BE_HOME/rms/bin/RMS.cdd`
6. Restart the RMS server.

RMS Server Configuration Property Reference

See [Configuring RMS Server Properties on page 569](#) for the related procedure.

Properties not documented in this table are either for internal use or are for other types of configuration explained elsewhere in the documentation.

Table 64 RMS Server Configuration Properties

Property	Notes
WS-RMS-Common Property Group	
<code>rms.checkin.revisionId.initValue</code>	<p>Initial value of the revision number to use for project check ins.</p> <p>Default is 10000</p>
<code>ws.artifact.deploy.location</code>	<p>Name of the directory for storing EAR files generated for deployment. Separate directory for each project is created, where EAR files for the project are stored.</p> <p>See also <code>be.codegen.rootDirectory</code> below.</p> <p>Default is <code>BE_HOME/rms/shared</code></p>
<code>be.codegen.rootDirectory</code>	<p>Name of the directory for storing class files generated by the Generate Deployable menu option. The directory location is relative to the <code>ws.artifact.deploy.location</code> directory.</p> <p>If this property is not present, files are stored in the deployment directory.</p> <p>If this property is present but has no value, a directory is created with the default name <code>Codegen</code>.</p>
<code>be.codegen.useLegacyCompilation</code>	<p>Set the value to true, if you want to use the file-based legacy compiler to build the EAR file. By default, the value is false, that means EAR files are built in memory.</p>

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
<code>ws.projects.customLib.location</code>	<p>Path of the directory for storing custom function libraries for a WebStudio project. The custom function libraries are saved in a folder (with the same name as of the project name) inside the directory identified by the <code>ws.projects.customLib.location</code> property.</p> <p>For example, if the value of the <code>ws.projects.customLib.location</code> property is <code>C:/Tibco/BE51/be/5.1/rms/custom-lib</code>, then the custom function libraries for the <code>MyTestProject</code> project should be saved in the location <code>C:/Tibco/BE51/be/5.1/rms/custom-lib/MyTestProject</code>.</p> <p>Default is <code>BE_HOME/rms/custom-lib</code>.</p>
RMS Property Group	
<code>rms.project.workspace</code>	<p>Not used in this release.</p> <p>Default is <code>workspace</code></p>
<code>rms.project.decisiondata</code>	<p>Not used in this release.</p> <p>Default is <code>decisiondata</code></p>
<code>rms.project.deployment</code>	<p>Not used in this release.</p> <p>Default is <code>deployment</code></p>
<code>rms.workflowstages.config.file</code>	<p>Not used in this release.</p> <p>Default is <code>BE_HOME/rms/config/RoleWorkflowStages.xml</code></p>
<code>rms.lockworkflowstages.config.file</code>	<p>Not used in this release.</p> <p>Default is <code>BE_HOME/rms/config/LockWorkflowStages.xml</code></p>

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
<code>rms.roleArtifactTypes.config.file</code>	<p>Not used in this release.</p> <p>Default is <code>BE_HOME/rms/config/RoleApplicableArtifactTypesConfig.xml</code></p>
<code>rms.external.entities.autodetect</code>	<p>Used internally. Do not change the value.</p>
<h3>Authentication and Authorization (ACL) Properties</h3> <p>User authentication topics are common to TIBCO BusinessEvents and the add-on products that use authentication. Common options are more fully documented in <i>TIBCO BusinessEvents Administration</i>. The relevant properties are also listed below for your convenience.</p>	
<code>be.auth.type</code>	<p>Specifies the authentication mechanism used, <code>file</code> or <code>ldap</code>. Additional configuration is required.</p> <p>Note Authentication topics are documented in the <i>TIBCO BusinessEvents Administration</i> guide.</p>
<code>be.auth.file.location</code>	<p>Location of the authentication file used for file-based authentication.</p> <p>Note Authentication topics are documented in the <i>TIBCO BusinessEvents Administration</i> guide.</p> <p>Default is <code>BE_HOME/rms/config/security/users.pwd</code></p>
<code>java.security.auth.login.config</code>	<p>Location of the JAAS login configuration file. You can substitute a different implementation of the JAAS login module than the one provided.</p> <p>Note Authentication topics are documented in the <i>TIBCO BusinessEvents Administration</i> guide.</p> <p>Default is <code>BE_HOME/rms/config/security/jaas-config.config</code>.</p>

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
<code>ws.projects.acl.location</code>	<p>Location of the directory used for all ACL (authorization) files. Files must be named using the format <i>RMSProjectName.ac</i>.</p> <p>Default is <i>BE_HOME/rms/config/security</i></p>
RMS-GVs Property Group	
<p>This group of properties override the values of global variables provided in the TIBCO BusinessEvents Studio BRMS project. They are made available as global variables so that they can be overridden at runtime, as needed.</p> <p>You can override the default global variable values in the <i>RMS.cdd</i> file (as shown below) for command-line startup. For details on overriding global variable values when deploying with TIBCO BusinessEvents Monitoring and Management or with TIBCO Administrator, see <i>TIBCO BusinessEvents Administration</i>.</p> <p>For defining and using global variables in TIBCO BusinessEvents Studio projects, see Working with Global Variables on page 14.</p>	
<code>tibco.clientVar.RMS/hostname</code>	<p>Specifies the host name or IP address of the machine where RMS is hosted. Remote clients can connect to the server at this location.</p> <p>Default is <i>localhost</i>.</p>
<code>tibco.clientVar.RMS/port</code>	<p>Specifies the port number of the machine where RMS hosted. This port is used for listening to client requests. See notes for <code>tibco.clientVar.RMS/hostname</code>.</p> <p>Default is <i>5000</i>.</p>
<code>tibco.clientVar.RMS/Approval/adminRole</code>	<p>The value specified here has Administrator role permissions. When you change the default value, the Administrator role still has Administrator permissions.</p> <p>Default is <i>Administrator</i>.</p>

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
WS Property Group	
<code>tibco.clientVar.Webstudio/warDir</code>	<p>Path to the Web archive for the client side representation of the TIBCO BusinessEvents WebStudio.</p> <p>Default is <code>rms/bin/WebStudio.war</code>.</p>
<code>tibco.clientVar.Webstudio/hostname</code>	<p>Specifies the host name or IP address of the machine where RMS server is hosted for WebStudio.</p> <p>Default is <code>localhost</code>.</p>
<code>tibco.clientVar.Webstudio/port</code>	<p>Specifies the port number of the machine where RMS server is hosted for WebStudio. This port is used for listening to client requests. See notes for <code>tibco.clientVar.RMS/hostname</code>.</p> <p>Default is <code>8090</code>.</p>
<code>tibco.clientVar.Webstudio/Connection/timeout</code>	<p>Timeout interval for the HTTP channel. This value is in milliseconds.</p> <p>Default is <code>60000</code>.</p>
<code>tibco.clientVar.Webstudio/sessionTimeOut</code>	<p>Time interval for which an inactive user session will be valid. This value is in minutes.</p> <p>Default is <code>30</code>.</p>
<code>ws.validatedT.temp.dir</code>	<p>Directory which is used as an interim storage of Decision Table contents during the validation process.</p> <p>Default is <code>rms/temp-dir</code>.</p>

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
Global Variable Properties for One-Way SSL Between RMS Server and Decision Manager Clients	
See RMS Server Configuration Property Reference on page 575 for the procedure.	
<code>tibco.clientVar.RMS/security/securePort</code>	The port to be used for secure HTTP communication. Default is 443.
<code>tibco.clientVar.RMS/security/sslCertificateStore</code>	Absolute path of the JKS keystore which contains the RMS server's public/private keys. If the certificate used is CA signed, include the CA's certificate.
<code>tibco.clientVar.RMS/security/sslCertificateStorePassword</code>	The password for the keystore.
Properties for One-Way SSL Between LDAP Server, RMS Server, and Decision Manager Clients	
Note	
<ul style="list-style-type: none"> • For one-way SSL between LDAP server and RMS server, you must also change the value of <code>be.auth.ldap.port</code>. • For one-way SSL between RMS server and Decision Manager clients, you must also set the value of <code>be.auth.ldap.port</code>. 	
See RMS Server Configuration Property Reference on page 575 for the procedure.	
<code>be.auth.ldap.usessl</code>	<p>If set to true, enables SSL between the RMS server and the LDAP server</p> <p>If set to false or not set, the LDAP authentication uses a non-secure channel.</p>
<code>javax.net.ssl.trustStore</code>	The absolute path of the keystore containing the LDAP server certificate chain.
<code>javax.net.ssl.trustStorePassword</code>	The password for the keystore.

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
<code>javax.net.ssl.trustStoreType</code>	The keystore type for keystores other than JKS type, for example, PKCS12.
Email Notification Property Group	
Email notification properties are used to send email notification when an artifact (business rule / decision table) undergoes a status change in approval workflow, such as, on artifact commit and on artifact reviews (Approve, Reject, BuildAndDeploy).	
In addition to these properties, you may have to define any SMTP protocol specific properties (<code>mail.smtps.auth</code> , <code>mail.smtps.ssl.enable</code>), depending on the mail server configuration. These properties should be defined with a prefix <code>ws.notify.prop</code> , for example, <code>ws.notify.prop.mail.smtps.auth</code> , <code>ws.notify.prop.mail.smtps.ssl.enable</code> .	
<code>ws.notify.enabled</code>	<p>If set to true, send email notification.</p> <p>If set to false, do not send email notification.</p> <p>Default is false.</p>
<code>ws.notify.mail.domain</code>	Specifies the domain on the email. This is used with RMS username to construct the email address to which notification is to be send. For example, if RMS username is "admin" and the property is set to "tibco.com", the email address for notification is admin@tibco.com.
<code>ws.notify.prop.MAIL_PROTOCOL</code>	<p>Specifies the email protocol to be used to send email. SMTP and SMTPS are valid values.</p> <p>Default is SMTPS.</p>
<code>ws.notify.prop.MAIL_SERVER_HOST</code>	Specifies the host name of the email server.
<code>ws.notify.prop.MAIL_SERVER_PORT</code>	Specifies the port at which mail server is listening.

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
<code>ws.notify.prop.SENDER_EMAIL</code>	Specifies email address to be used for sending the notification email.
<code>ws.notify.prop.SENDER_USERNAME</code>	Specifies the sender's username for authentication to email server.
<code>ws.notify.prop.SENDER_PASSWORD</code>	Specifies the sender's email password for authentication to email server.
Backing Store Property Group	
Backing Store property group enables you to add a shared JDBC connection without the need to build the new EAR file (<code>RMS.ear</code>) for the BRMS project.	
<code>tibco.clientVar.Webstudio/DB/driverName</code>	The name of the JDBC driver class. Default value (as shipped) is <code>oracle.jdbc.OracleDriver</code> .
<code>tibco.clientVar.Webstudio/DB/url</code>	The URL to use to connect to the database. You must supply the portions of the URL, for example, the host, port number, and database instance name. . Default value (as shipped) is <code>jdbc:oracle:thin:@localhost:1521:orcl</code>
<code>tibco.clientVar.Webstudio/DB/maxConnections</code>	The maximum number of database connections to allocate. The default maximum is 5. The minimum value that can be specified is 1.
<code>tibco.clientVar.Webstudio/DB/userName</code>	User name to use when connecting to the database.
<code>tibco.clientVar.Webstudio/DB/password</code>	Password to use when connecting to the database.

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
<code>tibco.clientVar.Webstudio/DB/timeout</code>	<p data-bbox="385 291 1330 418">Time (in seconds) to wait for a successful database connection. Only JDBC drivers that support connection timeouts can use this configuration field. If the JDBC driver does not support connection timeouts, the value of this field is ignored. Most JDBC drivers support connection timeouts.</p> <p data-bbox="385 435 1330 470">Default value (as shipped) is zero, which means no timeouts.</p>
Source Control System (SCS) Property Group	
SCS property group enables you to checkout a WebStudio project resources from a SVN location and commit the updates.	
<code>ws.scs.rootURL</code>	<p data-bbox="385 685 1330 755">Specifies the location of the directory which contains the repository of RMS projects. This is the root location for RMS project directories.</p> <p data-bbox="385 772 1330 864">If repository is present in SVN, then provide the SVN URL of the repository of RMS projects as the value of <code>ws.scs.rootURL</code> and "svn" as the value of the <code>ws.scs.impl.type</code> property.</p> <p data-bbox="385 881 1330 916">Default value (as shipped) is <code>BE_HOME/examples/standard/WebStudio</code></p>
<code>ws.scs.impl.type</code>	<p data-bbox="385 989 1330 1024">Identifies the type of repository being referred. The values are:</p> <ul data-bbox="385 1041 1330 1119" style="list-style-type: none"> <li data-bbox="385 1041 1330 1076">• file (default) <li data-bbox="385 1093 1330 1119">• svn
<code>ws.scs.repo.dir</code>	<p data-bbox="385 1197 1330 1267">The local repository directory on your system. All checked out project resources are downloaded at this location.</p> <p data-bbox="385 1284 1330 1319">Default value is <code>BE_HOME/rms/repo</code>.</p>
<code>ws.scs.command.path</code>	<p data-bbox="385 1388 1330 1489">Absolute path to the SVN command-line client. RMS server uses the SVN command-line tool to connect to the SVN repository. Any SVN command line tool can be used, one commonly used one being Tortoise SVN CLI.</p>

Table 64 RMS Server Configuration Properties (Cont'd)

Property	Notes
<code>ws.scs.default.username</code>	Default username to connect to the SVN repository mentioned in the <code>ws.scs.rootURL</code> property.
<code>ws.scs.default.password</code>	<p>Default password, encoded into the Base64 format, to connect to the SVN repository mentioned in the <code>ws.scs.rootURL</code> property. For example, <code>Password123</code> should be encoded with Base64 format and entered as <code>UGFzc3dvcnQxMjM=</code>.</p> <p>Note WebStudio uses default username and password (after decoding from Base64 format) to check out and check in the objects from the svn repository. You also have the option to enter your SCS username and password (in the Preference portlet under the Settings tab) in the WebStudio UI which overrides the default username and password. Enter the actual password in the WebStudio UI and not a Base64 format encoded one.</p>

Chapter 33 **Working With Projects in WebStudio**

This chapter explains how to work with projects, for example, creating projects and committing their resources for approval.

Topics

- [Starting and Logging in to RMS, page 586](#)
- [Checking Out a Project, page 588](#)
- [Updating \(Synchronizing\) a Project, page 590](#)
- [Committing Artifacts for Approval, page 591](#)

Starting and Logging in to RMS

To start RMS, do one of the following:

- Execute `BE_HOME/rms/bin/be-rms.exe` (or `be-rms.sh`, depending on your operating system), with valid arguments. For example, on Windows you would open a command window in `BE_HOME/rms/bin` and execute:


```
be-rms.exe
```
- (Windows only) Select Start > All Programs > TIBCO > `TIBCO_HOME` > TIBCO BusinessEvents *version* > Start Rules Management Server.

The server is ready when you see a message in the command window like this:

```
Info [HTTP-Channel-Startup] - [driver.http] Channel server for HTTP
Channel [Port:5000] successfully started
```

To Log in to RMS from WebStudio

1. Type the URL `http://host:port/WebStudio/` in a web browser. Default host is localhost and default port is 8090. Host and port are configurable using properties in `RMS.cdd`. Make sure that the RMS is running before logging to TIBCO BusinessEvents WebStudio.
2. Type the username and password for TIBCO BusinessEvents WebStudio.
3. Click **Login** to sign in to RMS.

The application displays the TIBCO BusinessEvents WebStudio dashboard.



The system restricts multiple simultaneous login of the same username.

To Log out of RMS

In TIBCO BusinessEvents WebStudio, click **Sign out** to log out. In case of session timeout, server logs out user automatically.

Status Check for RMS Connection

Once TIBCO BusinessEvents Studio or a TIBCO BusinessEvents Decision Manager or TIBCO BusinessEvents WebStudio client log in to RMS, a background activity on the dashboard page periodically polls the RMS server to refresh Worklist items. You can configure the interval (in seconds) of these status checks by setting the following property in `studio.tra`:

```
rms.heartbeat.delay
```

By default, the delay is of two seconds.

If the server goes down, the menu options for RMS are automatically disabled. However, if you are already performing an RMS operation and the server goes down, the operation fails. The login option is enabled when the server comes up again.

Checking Out a Project

If RMS is used in your environment, then you must log into RMS and check out a project from RMS before you can work with the project resources in the TIBCO BusinessEvents WebStudio UI.

When generating an EAR file using the Generate Deployable option, the project resources from the RMS repository are used and not from the workspace.

Depending on your role or roles, you may not have permission to check out all resources or to do all of the documented tasks.

For instructions on updating a project you already checked out, see [Updating \(Synchronizing\) a Project on page 590](#).



The checkout dialog do not list the imported 4.x release projects, if the project name contains a dot(.). These imported projects should be renamed inside TIBCO BusinessEvents Studio to remove dot(.) from the project name.

To Check Out a Project in WebStudio

1. Open the checkout window using any of the following methods:
 - In the **Dashboard** tab, click the link in the **My Projects** section if no projects are checked out.
 - In the **My Workspace** tab, click **WebStudio**, with the down-pointing triangle in front of it, and select **RMS > Checkout**.
 - In the **My Workspace** tab, right-click inside the **Group Contents** section and select **RMS > Checkout**.
 - In the **My Workspace** tab, click the checkout icon  in the RMS toolbar.
2. In the checkout window, select the project you want to checkout from the **Project List** drop-down.
3. Click on the checkboxes of the files you want to check out and click **OK**. Alternatively, click on the header checkbox to select all the files for checkout and click **OK**.

In TIBCO BusinessEvents WebStudio the checked out project is displayed under Group Contents and the project files are loaded to your workspace. The Group Contents section displays the artifacts based on the groups (Projects/Business Rules/Decision Tables) selected in the Groups section. The Business Rules and Decision Tables group are selected by default after you check out a project.

The TIBCO BusinessEvents WebStudio displays the project files in the Group Contents section either as a list or in a tree structure. In the Group Contents section click on the project view icon to select the **Show items as tree** or **Show items as list** options to toggle between the tree view () and the list view (). List view is the default view.

Updating (Synchronizing) a Project

This section explains how to update the projects on your workspace with the most recent changes from the corresponding RMS projects.

It is recommended that you log into the RMS server, and update (synchronize) your local copy of a project before continuing to work on it. Other users may have checked in changes, additions, and deletions that are now available for checkout. There may be other changes, too.

You can choose what updates to accept. Changes and additions to individual decision tables and business rules are listed.

To Update a Project in WebStudio

1. In the Group Contents section, select the artifact you need to synchronize.



Synchronize command does not support multiple artifact selection using Shift or Ctrl key. Alternatively, select a folder/project to select all artifacts under that folder/project, when project files are displayed in a tree structure.

2. Open the synchronize window using any of the following methods:
 - Click **WebStudio**, with the down pointing triangle in front of it, and select **RMS > Synchronize**.
 - Right-click on the selected artifact and select **RMS > Synchronize**.
 - Click the synchronize icon  in the RMS toolbar.
3. Click on the checkboxes of the files you want to synchronize and click **OK**. Alternatively, click the header checkbox to synchronize all listed project files and click **OK**.

The selected artifacts are updated with the latest project files from RMS.

Committing Artifacts for Approval

After you are done with the additions, changes, or deletions to a project you checked out of RMS, you commit the artifacts for approval. You can make a commit request for a selection of artifacts, or all.

Copies of submitted artifacts are saved in cache, with the names of the users who submitted them, version numbers, and the status `Committed`. A task is created for the Approver role to review the check-in requests. (see [Chapter 36, Working with the Approval Workflow in RMS, on page 625](#)).

A Revision ID is assigned to each commit request. This ID is used in other dialogs to allow you to view details on the request.

If the project is checked out from the SVN repository, then similar to the file mode the committed artifacts are saved in the local cache till the approval. After the approval, the committed artifacts from the cache are saved to the local repository, identified by the value of the `ws.scs.repo.dir` property in `RMS.cdd` (see [Configuring RMS Server Properties on page 569](#)), and checked in to the SVN repository as well. The comments entered for commit are used as checkin comments for the SVN repository, after approval.

Viewing the Audit Trail



This is available only if TIBCO BusinessEvents Decision Manager add-on is installed.

You can view the entire audit record of previous status changes for a checkin. This is especially useful for multi-stage approval process so that different reviewers actually share views on why a particular checkin was approved or rejected.

To Commit Project Resources for Approval from WebStudio

1. In the **Group Contents** section, select the artifact you need to commit.



Commit command does not support multiple artifact selection using Shift or Ctrl key. Alternatively, select a folder/project to select all artifacts under that folder/project, when project files are displayed in a tree structure.

2. Open the commit window using any of the following methods:
 - Click **WebStudio**, with the down pointing triangle in front of it, and select **RMS > Commit**.
 - Right-click on the selected artifact and select **RMS > Commit**.
 - Click the commit icon  in the RMS toolbar.
3. In the Commit window enter the revision description in the **Comments** text box.

Comments field is a mandatory field and system will not commit the changes to the repository if the Comments field is empty.
4. Click on the checkboxes for the files you want to commit to RMS project and click **OK**. Alternatively, click the header checkbox to select all listed project files.
5. Click **Finish** after the system displays checkin successful message with the revision ID.

This chapter explains how to work with decision tables, in TIBCO BusinessEvents WebStudio.



Decision tables are available in TIBCO BusinessEvents WebStudio in edit mode only if you have TIBCO BusinessEvents Decision Manager add-on installed on your system.



If you are creating/modifying the decision table for the first time, then generate the EAR files before creating/modifying decision tables. See [Generating Deployable Files \(EAR and Class Files\) on page 632](#) on how to generate the EAR files for the project.

Topics

- [Working With Decision Tables, page 596](#)
- [Understanding Columns and Rows \(Rules\) in a Decision Table, page 600](#)
- [Working With Rows and Columns in WebStudio, page 602](#)
- [Supported Operators, page 609](#)
- [Analyze a Decision Table, page 611](#)
- [Validate a Decision Table, page 613](#)

Working With Decision Tables

This section presents process to perform various decision table operations in TIBCO BusinessEvents WebStudio.

Create/Remove a Decision Table

A decision table can only be created using a virtual rule function (VRF). For instructions on adding a VRF, see [Chapter 16, Rules and Rule Functions, on page 251](#).

You can delete a newly added table as well as any existing table. If the table you want to delete has been checked into the RMS project, then you must submit the deletion for approval. After approval, the deleted table is removed from the RMS project.

To Add a Decision Table to a Project

1. Select **Projects** group.
2. In the **Group Contents** section, right click the virtual rule function (VRF) you want to use and click **New Decision Table**.
3. Enter the decision table name in the Decision Table pop-up window and click **OK**.

New empty decision table is displayed in the work area.

To Remove a Decision Table

1. Do one of the following:
 - In **Group Contents**, right-click the decision table name and select **Delete**.
 - In **Group Contents**, select the decision table name. Then click the delete icon  in the **Edit** toolbar.
2. If the table has been committed to the RMS project, choose one of the following options as appropriate:
 - Commit the project or the parent folder so that the RMS master copy of the table is also deleted. OR
 - Synchronize the project to replace the deleted copy with the RMS project version of the table.

Edit a Decision Table

Open a decision table

In the **Group Contents** section, double click the decision table you want to open. You can now edit, analyze, and export the open decision table.

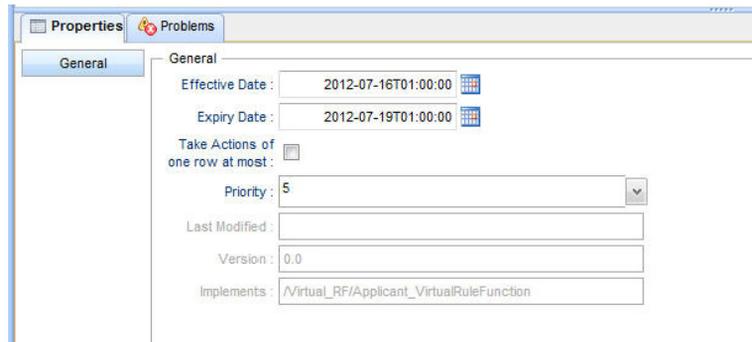
Add Rows and Columns

See [Working With Rows and Columns in WebStudio on page 602](#) for more details on adding rows and columns to a decision table in TIBCO BusinessEvents WebStudio.

Set Decision Table Properties

Follow below steps to set properties of a new decision table:

1. In the decision table work area, click the properties icon . The Properties view is displayed at the bottom of the work area.



2. In the **Effective Date** field, you can optionally define the date and time on which the decision table becomes valid in the runtime application. Click the calendar icon and use the Calendar dialog to set the date and then edit the time as needed.
3. Similarly in the **Expiry Date** field, you can optionally define the date and time after which the decision table is no longer valid in the runtime application.
4. Check the **Take Actions of One Row at Most** check box if you want the decision table to stop after one row (rule) passes the condition tests. Only the actions of that one row are taken (until the next time the decision table's rule function is called).

5. In the **Priority** field, set a priority as desired. When a VRF has multiple implementations (decision tables), the decision table's Priority setting determines the order in which the decision tables executes.
6. Click the Save icon  in the Edit toolbar to save the changes to the decision table.

The decision table is still open for you to edit and changes are not applied to RMS unless you commit the changes to RMS.

Export a Decision Table

You can export decision tables from TIBCO BusinessEvents WebStudio to Microsoft Excel spreadsheets. Sample exported Microsoft Excel Sheet:

Version										
	A	B	C	D	E	F	G	H	I	J
1	Version	2								
2										
3	Declarations									
4										
5	Path	Alias	Property							
6	/Concepts/applicant		BOTH							
7	/Events/Applicant		BOTH							
8										
9	Decision Table									
10										
11	Id	Condition	Condition	CustomA	Action (5)	Action (11)	Action (9)	Descriptio	Priority	
12		applicant.	applicant.	Custom A	applicant.	applicant.	applicant.	applicant.	Status	
13	1	"John Doe" >=20 && < System.de	true		2500			"VISA Granted"		
14	2	"Sandra W" >=20 && < System.de	false		0			"Loan Rejected"		
15	3	"Prakash" >=20 && < System.de	true		7500			"Pending"		
16	4	"Jane Doe" > 30	System.de	true	25000			"Platinum Status"		
17										
18	ExceptionTable									
19										
20	Id	Descriptio	Priority							
21										
22										

To Export the Decision Table

1. Open the decision table you want to export.
2. Execute the export command in either of the following ways:
 - Click the export icon  in the **Tools** toolbar. You see the **Save As** dialog for the exported file.
 - Right click the decision table in the Group Contents section and select **Export**.

3. Navigate to the directory in which you want to save the excel file and provide a file name.
4. Click **Save**.

Import a Decision Table

You can import correctly formatted Microsoft Excel file to create a decision table in TIBCO BusinessEvents WebStudio.



Only import from Excel files is supported. You cannot import from CSV (comma-separated values) files or tab-delimited text files. You can open such files in Excel and save as Excel format binary files.

To Import the Decision Table

1. Open the **Import** dialog in either of the following ways:
 - Select the virtual rule function and click the import icon  in the **Tools** toolbar.
 - Right click the virtual rule function in the Group Contents section and select **Import**.
2. Click **Choose File** in the **Import** dialog.
3. Browse the excel file you want to import and click **Open**.
4. Enter the new decision table name and click **OK**.

Understanding Columns and Rows (Rules) in a Decision Table

A decision table rule is a row in the decision table. It has one or more conditions and one or more actions. Each condition cell is equivalent to one condition (one line) in the TIBCO BusinessEvents rule editor condition area. Similarly each action cell is equivalent to one action (one line) in the TIBCO BusinessEvents rule editor action area (that is, one rule). Decision table rules are like business rules. The rule that calls the virtual rule function that implements the decision table participates in inferencing in the usual way. When the virtual rule function is called, the decision table rules are applied.

Conditions and Actions

The columns of a decision table are made up of condition columns on the left, and action columns on the right. Each column represents one condition or one action.

A condition is a test that must evaluate to true before the action is executed. If a decision table rule uses multiple conditions, all conditions for a row must evaluate to true in order for the action to execute.

In each row (rule) you define the specific conditions and actions. For example, if a condition column is Age (using a concept property of that name), then each row can define a different age range. The action for each row would define what action to take if a given concept instance contains an age property within the specified range.

If you add a second condition column called Income, then before the action is taken, a concept instance would be tested to see if both the age and the income are within the ranges specified in the rule's conditions.



Conditions that are blank, contain an asterisk, or are disabled are ignored, and are treated as if they evaluate to true.

Regular and Custom Conditions and Actions

A regular condition is the value of an entity specified in the VRF scope, or a simple comparison with the value - is greater than, is less than, is greater than or equal to, is less than or equal to.

A custom condition can use the rule language, standard functions, and data in the scope of the function at runtime (for example, scorecards and global variables). It can contain complex formulas.

A regular action sets the value of an entity specified in the VRF scope.

A **custom action** can use the rule language, standard functions, and data in the scope of the function at runtime to do whatever is desired. For example, the action could be to send an event out to a different system for follow-up.



Using non-literal values in a regular condition or regular action Suppose event A and event B are in scope, and event A has property ZZ, and event B has property YY. Both properties belong to the same data type. Suppose you then drag property ZZ to a condition or action column. In the cell, you can then specify a value as `b.YY`. The effect is different depending on the type of column:

- In a condition column, this means: compare the value of property ZZ with the value of property YY.
- In an action column this means: set the property ZZ to the value of YY.

Working With Rows and Columns in WebStudio

This section presents various operations you can perform as needed while working with rows and columns of a decision table in TIBCO BusinessEvents WebStudio.

Add Condition and Action Columns

The columns of a decision table are made up of condition columns on the left, and action columns on the right. Each column represents one condition or one action.

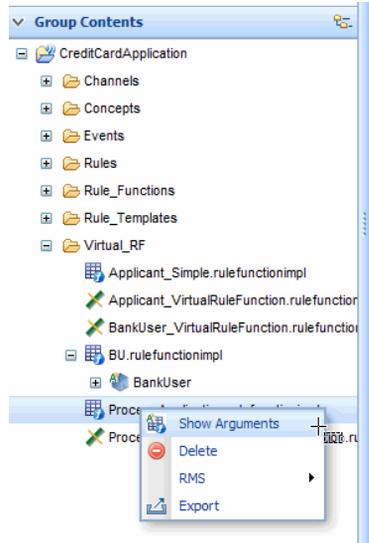
A condition is a test that must evaluate to true before the action is executed. If a decision table rule uses multiple conditions, all conditions for a row must evaluate to true in order for the action to execute.

Add condition and action columns from arguments under the VRF scope or using the rule language, standard functions, and data in the scope of the function at runtime (for example, scorecards and global variables).

To Add Condition and Action Columns from Arguments

1. Open the decision table you want to edit.
2. In the Group Contents section, right click the decision table and select **Show Arguments**.

A plus icon is displayed in front of the decision table for argument list expansion.

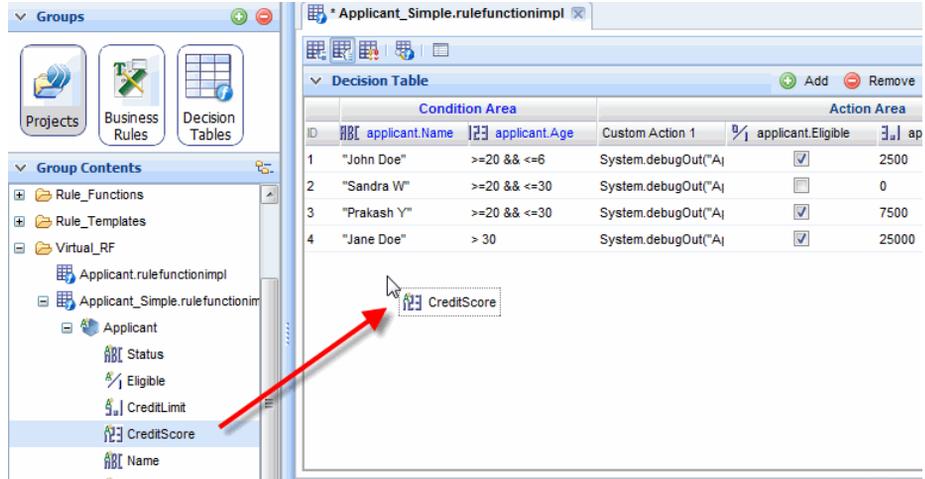


3. Click plus icon to expand the arguments list. If needed, click plus icons for concepts and rules as well.

The decision table arguments are listed under the decision table.

4. Drag a property from the argument list to the decision table editor.

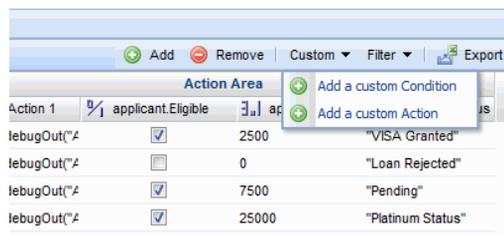
A Create Column pop-up window appears to create condition or action column.



- Click on the Condition or Action radio button and click **OK** to create a Condition or Action column. New Condition or Action column is added at the end of the Condition or Action column-sets respectively.

To Add Custom Condition and Action Columns

- Open the decision table you want to edit.
- Click **Custom** in the decision table editor menu and select either of the following options:
 - Select **Add a custom Condition** to insert a custom condition column at the end of condition column set.
 - Select **Add a custom Action** to insert a custom action column at the end of action column set.



Working with Columns

You can perform actions on a column to resize, sort, associate to domain model, or remove from a decision table.

Resize a column

Perform any of the following actions to resize a column in the TIBCO BusinessEvents WebStudio:

- Drag the right boundary of a column to manually resize them.
- Open the column header menu using either of the following ways and click **Auto Fit** to resize the column to fit to the content size:
 - Right click on the column header you want to resize.
 - Move the mouse pointer to the column header to view the down-pointing triangle and click the triangle.

Decision Table		
ID	applicant.Age	applicant.CreditScore
1		
2		0 && < 50000
3		= 500 && < 600
4		= 500 && < 600
5		= 500 && < 600
6		
7		= 600 && < 700
8		= 600 && < 700
9	>= 16 && < 18	>= 700 && <= 850

- Double click on the right boundary of a column to make it auto fit to the content size.
- Open the column header menu using either of the above mentioned ways and click **Auto Fit All Columns** to resize all columns to fit their content size.

Sort the decision table

Sort the decision tables in TIBCO BusinessEvents WebStudio in one of the following ways:

- Click once on the column header to sort the table in ascending order of the column and click again on the column header to re-sort the table in descending order of the column.

- Open the column header menu using either of the following ways and click **Sort Ascending** or **Sort Descending** to sort the table in ascending or descending order of the column:
 - Right click on the column header you want use for sorting.
 - Move the mouse pointer to the column header to view the down-pointing triangle and click the triangle.



Sort options are not available for columns with boolean entry values and for custom condition and action columns.

Associate and Remove Domain Association

A domain model specifies the values that you may find useful for defining ontology item properties. For example, instead of typing text for a certain concept property, you can pick a value from a list, or enter a value within a predefined range.

To associate a domain model, right click on the column and from the menu click **Associate Domain**. The setting toggles to ON. To clear the domain association click **Associate Domain** again. The setting toggles to OFF.

Remove a Column

Right click the column header and click **Remove** to remove the column from the decision table.



You cannot remove the last condition column because a minimum of one condition column must exist if you have any action columns.

Add a Row

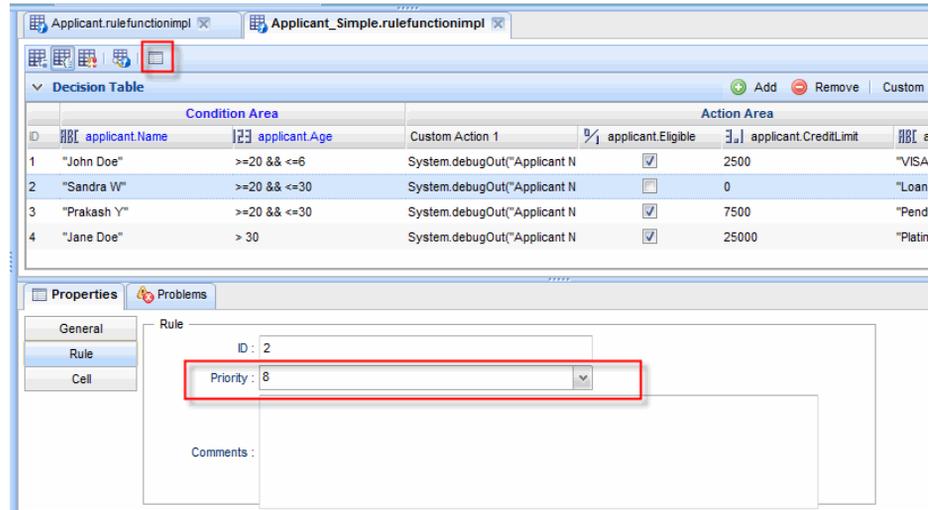
Each row in a decision table represents a separate business rule. You can control the order in which sets of rows are executed using the row priority setting, see [Setting Rule \(Row\) Priorities on page 607](#).

Set Row Priority in TIBCO BusinessEvents WebStudio

To set the row priority, with ten as the lowest priority and one as the highest, do the following:

1. Select the row whose priority you want to set.
2. Click the Properties icon  to display the Properties view.

3. Select the Properties view, and select the **Rule** side-tab.
4. In the **Priority** field, select the appropriate value.
5. Click the Save icon  in the Edit toolbar to save the changes to the decision table.



Setting Rule (Row) Priorities

Each row is like a separate business rule. You can control the order in which sets of rows are executed using the row priority setting. Rows with higher priorities are executed before those with lower priorities, as follows:

1. First, all conditions are checked for all rows that have the highest priority. (The checking order within a set of rows with the same priority is not determinate.)
2. Then the rule actions for all of those rows whose conditions evaluate to true are executed. (The execution order is not determinate. The runtime engine optimizes rule execution.)
3. The process is repeated for all rows with the next highest priority, and so on.

To Set Row Priority

Ten is the lowest priority and one is the highest. Five is the default priority.

1. Select the row whose priority you want to set.
2. Select the Properties view, and select the Rule side-tab.

- In the Priority field, select the appropriate value and click Apply.

Filter the Rows

Filtering rows helps you to focus on certain rows instead of entire table. Apply filters on multiple columns for more refined results. In TIBCO BusinessEvents WebStudio decision table filter makes only textual comparison between the row values and filter values.

Click **Filter > Decision Table** in the decision table editor menu to enable the filter for the decision table. Enter the values in filter cells (above column header) to filter the table.

To remove the filter click **Filter > Remove** in the decision table editor.

Decision Table						
Condition Area			Action Area			
ID	applicant.Name	applicant.Age	Custom Action 1	applicant.Eligible	applicant.CreditLim	applicant.Status
1	"John Doe"	>=20 && <=6	System.debugOut("Ap	<input checked="" type="checkbox"/>	2500	"VISA Granted"
4	"Jane Doe"	> 30	System.debugOut("Ap	<input checked="" type="checkbox"/>	25000	"Platinum Status"

Remove a Row

Select the row you need to remove and click **Remove** in the decision table editor. Click **OK** in the Remove Decision Table Record window to confirm removal of the row. You can remove only one row at a time.

Supported Operators

Table 65 lists the operators supported by Decision Tables:

Table 65 Supported Operators

Operator	Description
==	Equals to
!=	Does not equal to
>	Greater than
>=	Greater than or equals to
<	Less than
<=	Less than or equals to
	Logical OR
&&	Logical AND
+	Addition or String concatenation (depending on context)
!	Logical NOT (only with boolean)
-	Subtraction
*	“Don’t Care” (that is, ignore), or multiplication (depending on context)
/	Division
=	Assignment
.	Scope resolution
()	Operator precedence order or function call (depending on context)
[]	Array declaration or array indexing (depending on context)
{}	Block resolution or array Initialization (depending on context)
?:	Ternary conditional
++	Increment (used only with custom actions)

Table 65 Supported Operators (Cont'd)

Operator	Description
--	Decrement (used only with custom actions)

Using Operators in Tables

Named or Scoped Conditions must Evaluate to Boolean

For logical operators, named conditions will have to use them explicitly from the second operand onwards.

For example, if the LHS of the condition is the column name, `bankuser.Age`, the following are valid entries:

```
50 || ==60 , 40 && ==80, 20 || >=80
```

Custom Conditions must evaluate to Boolean

Same as above, except that you must specify the entire Boolean expression in the cell, including LHS.

Actions or Custom Actions

Named actions perform assignment of the cell expression to the column name property.

Custom actions are treated as normal strings and can be anything that you would enter in the THEN section of rules (LHS included).

Analyze a Decision Table

Table Analyzer will create a sparse matrix data structure representing an optimized form of the decision table in memory. In the Table Analyzer you can set example condition values and perform various validation checks on the currently displayed decision table.



- Table analyzer supports the Not Equal To (!=) operator.

The table analyzer feature provides various tools that you can use to check the validity of your decision tables.

To analyze a decision in the TIBCO BusinessEvents WebStudio, do the following:

1. Open the decision table you want to analyze.
2. Click the Decision Table Analyzer icon to view Table Analyzer.
3. In the Table Analyzer section, controls are created dynamically for setting values and ranges for each condition in the table. Enter or select the example values for which you want to analyze the table.

For example, if the condition is a range, for instance “< 40” or “> 10 && < 100”, then the corresponding control in the Table Analyzer section is a Slider that spans the range from the minimum value to the maximum value. You can then select a range of values within the range.

4. Click the Analyze icon  to analyze the decision table based on the selected values.

A pop-up window displays the count of any issues. The Problems section at the bottom displays more information on issues. Click on the issue to highlight the row with the issue.

The screenshot displays the WebStudio interface for a decision table. The main window is titled "Applicant_Simple.rulefunctionimpl". It features a "Decision Table" section with a table of conditions and actions. Below this is a "Table Analyzer" section with a dropdown for "applicant.Name" (set to "John Doe") and a slider for "applicant.Age" (set to 22). At the bottom, the "Problems" section shows a list of warnings, including "Uncovered range for c /Virtual_RF/Applicant_CreditCardApplication/ 1" and "Uncovered range for c /Virtual_RF/Applicant_CreditCardApplication/ 2".

ID	Condition Area	Action Area
1	"John Doe" >=20 && <= System.deb	<input checked="" type="checkbox"/> 2500 "VISA Grant"
2	"Sandra W" >=20 && <= System.deb	<input type="checkbox"/> 0 "Loan Reject"
3	"Prakash Y" >=20 && <= System.deb	<input checked="" type="checkbox"/> 7500 "Pending"
4	"Jane Doe" > 30 System.deb	<input checked="" type="checkbox"/> 25000 "Platinum Sta"

Type	Description
Warnings (10)	
	Uncovered range for c /Virtual_RF/Applicant_CreditCardApplication/ 1
	Uncovered range for c /Virtual_RF/Applicant_CreditCardApplication/ 2

You can click the Show Coverage icon  in the table analyzer section to highlight all rows that meet the criteria.

For example, if you select Male for Gender; true for Eligible; and 50000-60000 for Income, clicking Show Coverage displays all rows that have Male, true, and an income between 50000 and 60000.

Validate a Decision Table

In TIBCO BusinessEvents WebStudio, you can validate decision tables for the changes using the validate command. Select the decision table and click on the validate icon  to validate the table.



The validate icon is activated when the decision table is opened in the editor. EAR file for the project should be present in the deploy location; otherwise, validate command gives error.

If there are any access control violations or syntax errors in the tables, they are shown in the Problems View tab at the bottom of the application. New errors in syntax are added to these existing errors. Double-click errors to see the problematic view. Take any needed corrective actions and then validate the table again until all errors are resolved.

Chapter 35 **Business Rules**

Business Rules artifact allows non-technical business users to build complex business rules in TIBCO BusinessEvents WebStudio. Business Rules are created using the Rule templates developed in TIBCO BusinessEvents Studio. See [Chapter 17, Rule Templates, on page 279](#) for details on how to create rule template. This chapter explains about business rules and how to create them.

Topics

- [Business Rules Overview, page 616](#)
- [Working with Business Rules, page 617](#)
- [Business Rule Operators, page 622](#)

Business Rules Overview

Business rules are defined inside of TIBCO BusinessEvents WebStudio, and are intended to be used by non-technical users. TIBCO BusinessEvents WebStudio users complete the definition of the rule instance using one of the following types of user interface:

- A form where users can provide values for bindings defined in TIBCO BusinessEvents Studio. The form interface is the visual presentation of the rule template defined by rule template views. See [Rule Template Views on page 283](#) for more details on rule template views.
- A builder style of user interface where users can build conditions and actions using various operators. A builder based user interface is presented to user only if there is no view associated with the rule template.

Business Rule Builder

The builder for business rule in the TIBCO BusinessEvents WebStudio consists of two sections:

When section - condition builder

When section lets user define additional conditions that must be met, in addition to the pre-condition defined in the Rule Template. The builder determines the scope of the rule template, and allows the user to define individual conditional clauses based on that scope. Each conditional clause has one or more sub-clause, and a match operator determines whether conditional clause is true when all sub-clause are true, or any one sub-clause is true, or none of the sub-clauses evaluate to true.

Then section - command builder

Then section defines additional actions to be taken based on the conditions defined. The builder lets user define actions based on the statements defined in the Action Context section of the rule template.



Ensure that all the actions in the business rule are defined and the order should match that of the rule template. Otherwise, the business rule will be invalid.

Working with Business Rules

In TIBCO BusinessEvents WebStudio, you can create, modify, or delete business rules. Any modification after commit goes through the RMS approval process (see [Chapter 36, Working with the Approval Workflow in RMS, on page 625](#)).

A business rule in a project is identified by the `.ruletemplateinstance` file extension. The rule template instance file is generated when `BuildAndDeploy` option is selected for the business rule in approvals. See [Generating One Business Rule's Rule Template Instance File on page 635](#) for more details.

Adding Business Rules

In TIBCO BusinessEvents WebStudio, business rule is created using a rule template defined in TIBCO BusinessEvents Studio. If a rule template view is associated for the rule template, the business rule has a easy-to-use form based interface to specify conditions and action; otherwise, condition and actions are added using builder. See [Rule Template and Rule Template View Overview on page 280](#) for overview on rule template, and rule template view.

See [Chapter 33, Working With Projects in WebStudio, on page 585](#) for more details on starting and logging to RMS, and checking out project resources.

To add a business rule

1. In TIBCO BusinessEvent WebStudio, under the Group Contents section, right click on the rule template you want to use and select **Create Business Rule**.
2. Type the business rule name in dialogue box and click **OK**.
3. If the rule template view is defined for the rule template then HTML form is presented in the business rule editor. Then follow steps in [To add a business rule using HTML form on page 617](#) to create a business rule.
4. Otherwise, if view is not associated with the rule template, a builder interface open up in the editor. Follow the steps in the following sections to create a business rule:
 - [To add a conditional clause using builder on page 618](#)
 - [To add a action using builder on page 619](#)

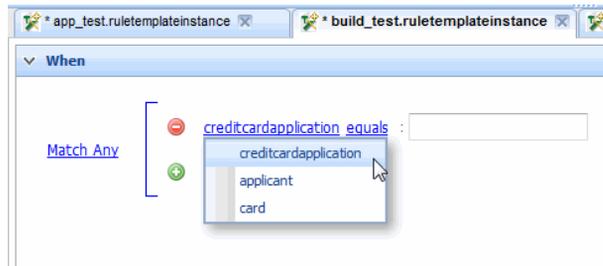
To add a business rule using HTML form

1. In WebStudio, enter the threshold values in the input fields of the HTML form.

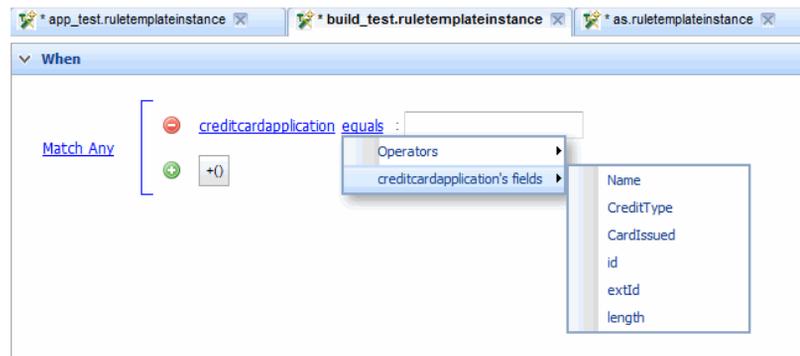
2. Click the save icon  to save updates to the business rule.

To add a conditional clause using builder

1. In the business rule editor, under the When section, click on artifact-name link and select the artifact you require.



2. Click on the operator-name link to select another operator, or a child artifact of the previously selected artifact, as per your requirement. See [Condition Builder Operators on page 622](#) for list of supported operators.

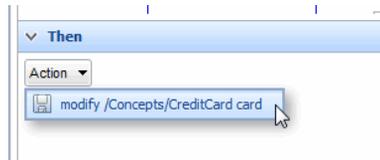


3. Enter the conditional value for the artifact in the empty box.
4. Click the add icon  to another condition to the conditional clause, or click the remove icon  to remove the existing condition from the conditional clause, if you want.
5. Click the add sub-clause icon  to add conditional sub-clause to the conditional clause, if you want.

6. Click match-operator link to select the matching condition for which the conditional clause evaluates to true. The values are:
 - **Match Any** Conditional clause is true if any of the condition evaluates to true
 - **Match All** Conditional clause is true only if all of the conditions evaluates to true
 - **Match None** Conditional clause is true only if none of the conditions evaluates to true.
7. Click the save icon  to save updates to the business rule.

To add a action using builder

1. In the business rule editor, under the Then section, click **Action** dropdown to select the actions available for the rule template.



2. Click the artifact-name link and select the artifact which you want modify.



3. Click the operator-name link and select the command for the artifact. See [Action Builder Operators on page 624](#) for list of available on operators.
4. Type the value, you want to set for the artifact, in the empty box.
5. Click the add icon  to add another parameter for the action, or click the remove icon  to remove the parameter for action, if you want.
6. Click the save icon  to save updates to the business rule.

Export a Business Rule

You can export business rules from TIBCO BusinessEvents WebStudio to rule template instance file. Sample exported rule template instance file content:

```
<?xml version="1.0" encoding="UTF-8"?>
<RuleTemplateInstance id="738cd4ea-8569-4744-aa86-fb88a89bf0a3"
implementsPath="/Rule_Templates/Applicant_PreScreen">
  <actions/>
  <binding id="minimumIncome" value="5000"/>
  <binding id="minimumAge" value="25"/>
  <binding id="creditType"/>
</RuleTemplateInstance>
```

To Export the Business Rule

1. Open the **Export** dialog in either of the following ways:
 - Select the business rule and click the export icon  in the **Tools** toolbar.
 - Right click the business rule in the Group Contents section and select **Export**.
2. Navigate to the directory in which you want to save the rule template instance file and provide a file name.
3. Click **Save**.

Deleting Business Rules

If the business rule you want to delete has been checked into the RMS project, then you must submit the deletion for approval. After approval, the deleted table is removed from the RMS project, but it is not removed from user's workspace.



If you are not sure whether a table you deleted was committed to the RMS project in a prior action, select Project > Commit. If the deleted table does not appear in the Changes Made panel, then it is a local table.

To Delete a Business Rule

1. Do one of the following:
 - In **Group Contents**, right-click the business rule name and select **Delete**.
 - In **Group Contents**, select the business rule name. Then click the delete icon  in the **Edit** toolbar.

2. If the business rule has been committed to the RMS project, choose one of the following options as appropriate:
 - Commit the project or the parent folder so that the RMS master copy of the business rule is also deleted (see [To Commit Project Resources for Approval from WebStudio on page 591](#)). OR
 - Synchronize the project to replace the deleted copy with the RMS project version of the business rule (see [To Update a Project in WebStudio on page 590](#)).
3. Do one of the following:
 - In Studio Explorer, right-click the business rule name and select **Delete**.
 - In Studio Explorer, select the business rule name. Then select **File > Delete**.
4. If the business rule has been committed to the RMS project, choose one of the following options as appropriate:
 - Select **Project > Commit** so that the RMS master copy of the table is also deleted. OR
 - Select **Project > Synchronize** to replace the local deleted copy with the RMS project version of the table.



When an artifact deletion is approved, the artifact is permanently deleted from the RMS project. However, When Decision Manager or WebStudio users update their projects, the deleted artifacts do not appear in the list of changes.

Users who have already checked out the project must manually delete their local copy of the artifact.

Business Rule Operators

This section lists the operators supported in the business rules builder. RTI builder operators consists of a conditional clause operators and action builder operators. The operator list changes as per the parent entity type in the builder.

Condition Builder Operators

Parent Type	Operator Value on WebStudio Client	Operator Value on RMS Server
<i>Concept / Event</i>		
	Matches Other Field	== <field>
	Differs From Field	!= <field>
	Is Null	== null
	Is not Null	!= null
<i>Data / Time / Datetime</i>		
	Matches Other Field	== <field>
	Differs From Field	!= <field>
	Is Null	== null
	Is not Null	!= null
<i>Integer / Float</i>		
	Matches Other Field	== <field>
	Differs From Field	!= <field>
	Is Null	== null
	Is not Null	!= null
	Greater Than	>
	Greater Than Field	> <field>

Parent Type	Operator Value on WebStudio Client	Operator Value on RMS Server
	Greater Than Equal To	>=
	Greater Than Equal To Field	>= <field>
	Less Than	<
	Less Than Field	< <field>
	Less Than Equal To	<=
	Less Than Equal To Field	<= <field>
	Equals	==
	Not Equals	!=
<i>Text / String</i>		
	Matches Other Field	== <field>
	Differs From Field	!= <field>
	Is Null	== null
	Is not Null	!= null
	Equals	==
	Not Equals	!=
<i>Boolean</i>		
	Equals	==
	Not Equals	!=
	Matches Other Field	== <field>
	Differs From Field	!= <field>

Action Builder Operators

Parent Type	Operator Value on WebStudio Client	Operator Value on RMS Server
<i>Concept / Event</i>		
	Set To	=
	Set To Field	= <field>
	Set To Null	= null
<i>Data / Time / Datetime</i>		
	Set To	=
	Set To Field	= <field>
	Set To Null	= null
<i>Integer / Float</i>		
	Set To	=
	Set To Field	= <field>
	Increment By	+=
	Increment By Field	+= <field>
	Decrement By	-=
	Decrement By Field	-= <field>
<i>Text / String</i>		
	Set To	=
	Set To Field	= <field>
	Set To Null	= null
<i>Boolean</i>		
	Set To True	= true
	Set To False	= false
	Set To Field	= <field>

Working with the Approval Workflow in RMS

This chapter explains the approval workflow of decision tables or business rules.

Topics

- [Working with the Approval Workflow—Overview, page 626](#)
- [Checking a Worklist and Taking Action, page 628](#)

Working with the Approval Workflow—Overview

When business users have finished working with artifacts (decision tables or business rules), they submit them for approval. All business user actions, such as committing or deleting artifacts are committed to the RMS project only when a user whose role gives permission to do so approves the requests.

The approval permissions are set at project level using the ACL file for the project. Allow the approval action type for the user role in the ACL file to grant approval permission. See *TIBCO BusinessEvents Administration* for details on setting permission for a user role at project level.

This section explains the provided workflow for approving and rejecting business users' commit requests.



This manual documents the behavior of the product as shipped. The workflow, the roles used, and the permissions granted to the roles are all configurable. However, the general flow is likely to be similar to that described in this guide.

With this release, the approval can only be done in TIBCO BusinessEvents WebStudio.

Use of RMS—and therefore the approval process—is required when TIBCO BusinessEvents WebStudio is used.

See [Committing Artifacts for Approval on page 591](#) for business user procedures.

See [Checking a Worklist and Taking Action on page 628](#) for approver procedures.

Who Can Approve or Reject Commit Requests

In TIBCO BusinessEvents WebStudio users with Administrator role permissions can approve or reject commit requests, or delegate approval. They can also delegate the task of approving or rejecting commit requests to any specified user role and not to a specific user. In this case, the approval or rejection requests will be sent to all the users in that role. See [To Delegate a Workitem on page 628](#).

Approval Status Values

COMMITTED When a user commits one or more resources for approval, the status of the request is set to Committed. The approver then sets the status as appropriate.

APPROVE The request was approved. If the submission request was to add or change an artifact, that artifact is copied to the RMS project. If the submission request was to delete an approved artifact, that artifact is deleted from the RMS project.

REJECT The change was rejected. If the submission request was to add or change an artifact, that artifact is not copied to the RMS project. If the submission request was to delete an approved artifact, that artifact is not deleted. It is up to the business user to make changes in their local project accordingly.

BUILDANDDEPLOY The artifact is build and ready to be deployed. The artifact after approval is built and kept ready to be deployed at the location specified by a property in the RMS .cdd file. See [RMS Server Configuration Property Reference on page 575](#) for more details on property.

Checking a Worklist and Taking Action

Approval requests are handled in the Worklist window in TIBCO BusinessEvents WebStudio.

To Check Your Worklist and Take Action

1. Open the worklist window using any of the following methods:
 - Click **WebStudio**, with the down pointing triangle in front of it, and select **RMS > Show Worklist**.
 - Right-click on any artifact and select **RMS > Show Worklist**.
 - Click the show worklist icon  in the RMS toolbar.

The worklist window displays all the requests submitted for approval.

2. To view the artifacts submitted for approval, click on the triangle in front of the Revision ID. To examine the artifact before taking an action, double-click the artifact. The decision table or business rule displays in the editor.
3. To take action on a artifact, click on the `Committed` status under the Status column, and select a value from status drop-down menu:
 - Approve
 - Reject

A message field displays error messages in the event of a problem.

4. Click anywhere outside the drop-down menu and click **Ok**.

To Delegate a Workitem

A user in the Administrator role can delegate a workitem to one or more roles. Users belonging to any of those roles see the workitem in their worklist.



Users in the role specified by the RMS global variable property `tibco.clientVar.RMS/Approval/adminRole` have Administrator role privileges. See [Configuring RMS Server Properties, page 569](#).

1. Open the worklist window using any of the following methods:
 - Click **WebStudio**, with the down pointing triangle in front of it, and select **RMS > Show Worklist**.
 - Right-click on any artifact and select **RMS > Show Worklist**.
 - Click the show worklist icon  in the RMS toolbar.

The worklist window displays all the requests submitted for approval.

2. Check the checkbox for one or more Revision IDs, you need to delegate.
3. Select one or more user roles from the **DelegateTo** dropdown and click **Apply**. You see all roles that you do not belong to, and that have permission to approve or reject the artifact.

The workitem disappears from your worklist. It can be seen in the worklist of the users in roles you selected in [step 3](#).

Deployable Files Generation and Deployment for WebStudio

This chapter explains configuration and procedures for generating EAR files, decision table classes, and business rule's rule template instance file for TIBCO BusinessEvents WebStudio.

Topics

- [Generating Deployable Files \(EAR and Class Files\), page 632](#)
- [Business Rule Hot Deployment, page 636](#)

Generating Deployable Files (EAR and Class Files)

Deployable files are generated using the TIBCO BusinessEvents Studio project files, located in the RMS project's folder.

The deployable files are of two types: EAR files and class files.

EAR files are generated in Decision Manager and TIBCO BusinessEvents WebStudio UI.

You can generate class files using the Decision Manager and TIBCO BusinessEvents WebStudio UI, or at the command line. You can generate an entire project's class files, or one decision table's class file, or one business rule's rule template instance file. If you generate an entire project's class files, a property in the CDD file enables you to exclude unwanted packages.



You can generate one decision table's class file and business rule's rule template instance file using TIBCO BusinessEvents WebStudio only.

Location of Generated Files

Location of generated files can be configured using deploy location properties in the RMS.cdd file. See [RMS Server Configuration Property Reference on page 575](#) for more details on deploy location properties. Below are the location of generated files as per the preconfigured location in the shipped product.

- EAR files are saved in the RMS project directory in the RMS's shared directory. For example, EAR file for CreditCardApplication project (`CreditCardApplication.ear`) is stored at
`BE_HOME\rms\shared\CreditCardApplication`
- Project's class files are saved in the `codegen` subdirectory under the RMS project directory in the RMS's shared directory. For example, class files for CreditCardApplication project (`CreditCardApplication.ear`) are stored at
`BE_HOME\rms\shared\CreditCardApplicatio\codegen`
- Individual decision table's class file and business rule's rule template instance are saved in the RMS project directory in the RMS's shared directory.

Generating the Project EAR or All Project Class Files

You can generate deployable EAR files only if you have permission to check out all the required project resources (ACLs may limit what you can check out).



If you have created/modified the decision table for the first time and did not generate the EAR files before creating/modifying decision tables, then you receive error while generating the deployables:

```
EAR for projects: <project_name> either not present or not in sync
at deploy location.
```

To Generate the Project EAR or All Project Class Files in WebStudio

1. Open the Generate Deployables window using any of the following methods:
 - Click **WebStudio**, with the down pointing triangle in front of it, and select **RMS > Generate Deployables**.
 - Right-click on any artifact and select **RMS > Generate Deployables**.
 - Click the Generate Deployables icon  in the RMS toolbar.

The Generate Deployables window is displayed.
2. Do one of the following:
 - To generate class files, check the **Generate Classes Only** checkbox.
 - To generate EAR files, do not check the **Generate Classes Only** checkbox.
3. From the drop-down list, select the project for which you want to generate the deployable files.
4. (For EAR files only.) If you want to generate the debug information, check the **Generate Debug Info** checkbox.
5. (For EAR files only.) If you want to include all service-level global variables, check the checkbox.
6. Click **OK**.

Files are saved to the preconfigured location (see [Location of Generated Files on page 632](#)). Any existing files are overwritten.

Generating One Decision Table's Class File

In TIBCO BusinessEvents WebStudio, the BuildAndDeploy status menu option in the Worklist window generates the class file for just selected decision tables. You can generate a decision table's class file only after it is approved. See [Checking a Worklist and Taking Action on page 628](#) on how to approve a commit request.



You can generate one decision table's class file using TIBCO BusinessEvents WebStudio only.



If you have created/modified the decision table for the first time and did not generate the EAR files before creating/modifying decision tables, then you receive error while generating the deployables:

EAR for projects: <project_name> either not present or not in sync at deploy location.

To Generate One Decision Table Class File

1. Open the worklist window using any of the following methods:
 - Click **WebStudio**, with the down pointing triangle in front of it, and select **RMS > Show Worklist**.
 - Right-click on any artifact and select **RMS > Show Worklist**.
 - Click the show worklist icon  in the RMS toolbar.

The worklist window displays all the requests submitted for approval as well approved requests.

2. To view the approved decision tables, click on the triangle in front of the Revision ID. To examine the decision table before taking an action, double-click the table. The decision table displays in the editor.
3. To generate classes for the decision table, click on the *Approved* status under the Status column, and select **BuildAndDeploy** from the status drop-down menu. The project EAR must already be created.

A message field displays error messages in the event of a problem.

4. To start generating classes click anywhere outside the drop-down menu and click **Ok**.

Generating One Business Rule's Rule Template Instance File

In TIBCO BusinessEvents WebStudio, the BuildAndDeploy status menu option in the Worklist window generates the rule template instance file for just selected business rules. You can generate a business rule's rule template instance file only after it is approved. See [Checking a Worklist and Taking Action on page 628](#) on how to approve a committed request.

To Generate One Rule Template Instance File

1. Open the worklist window using any of the following methods:
 - Click **WebStudio**, with the down pointing triangle in front of it, and select **RMS > Show Worklist**.
 - Right-click on any artifact and select **RMS > Show Worklist**.
 - Click the **Show Worklist** icon  in the RMS toolbar.

The worklist window displays all the requests submitted for approval as well approved requests.

2. To view the approved business rules, click the triangle in front of the Revision ID. To examine the business rule before taking an action, double-click the rule. The business rule displays in the editor.
3. To generate rule template instance file for the business rule, click on the **Approved** status under the Status column, and select **BuildAndDeploy** from the status drop-down menu.

A message field displays error messages in the event of a problem.

4. To start generating rule template instance file, click anywhere outside the drop-down menu and click **Ok**.

Business Rule Hot Deployment

For hot deployment of a business rule, you need to have a generated EAR for the project (see [Generating the Project EAR or All Project Class Files on page 633](#)). Then, after you have updated the business rule and generated the rule template instance file (see [Generating One Business Rule's Rule Template Instance File on page 635](#)), you can hot deploy the classes to running TIBCO BusinessEvents engines using JMX.

To Hot Deploy Rule Template Instance using JMX

1. Open the CDD file of the project and add the property `be.cluster.ruletemplateinstances.deploy.dir` under the appropriate processing unit, for example `default`.

The property points to the location where artifacts are generated and identified by the property `ws.artifact.deploy.location` (see [WS-RMS-Common Property Group on page 575](#) for more details on the property). For example, the rule template instance file for the `CreditCardApplication` project is stored at

```
BE_HOME\rms\shared\CreditCardApplication
```

2. Save and close the CDD file.
3. Start the engine with the project EAR and CDD.
4. Open a command window in the bin directory of your JDK installation and type `jconsole` and press **Enter**. You see the New Connection dialog.
5. Select the connection corresponding to the TIBCO BusinessEvents node that you want to view (each node runs in one JVM) and click **Connect**. You see the multi-tabbed Connection dialog.
6. Select the **MBeans** tab and expand the tree nodes on the left to **com.tibco.be > Cluster > Operations**.
7. Specify the agent name (for example, `inference-class`) and click **loadAndDeployRuleTemplateInstances**.

A confirmation message is displayed that the rule template instance is deployed.

Now the rule template instance is in Rete, and on asserting events/concepts the appropriate rule template instance gets executed.

Chapter 38 **ActiveMatrix BusinessWorks Integration**

This chapter explains how to use the TIBCO ActiveMatrix BusinessWorks plug-in, to integrate ActiveMatrix BusinessWorks and TIBCO BusinessEvents applications.

Topics

- [Overview of Integration with ActiveMatrix BusinessWorks, page 638](#)
- [Integration Components, page 639](#)
- [Design Considerations, page 641](#)
- [Configuring a RuleServiceProvider Configuration Resource, page 642](#)
- [TIBCO BusinessEvents RuleServiceProvider Configuration Resource Reference, page 643](#)
- [Working With the TIBCO BusinessEvents Activities, page 645](#)
- [Receive Event Resource Reference, page 646](#)
- [Send Event Resource Reference, page 648](#)
- [Wait for Event Resource Reference, page 649](#)
- [Invoking a TIBCO BusinessEvents Rule Function from a Process, page 652](#)
- [Working With Invoke RuleFunction Activities, page 654](#)
- [Invoke RuleFunction Resource Reference, page 656](#)
- [Working with the BusinessWorks Functions, page 658](#)

Overview of Integration with ActiveMatrix BusinessWorks

If TIBCO ActiveMatrix BusinessWorks used in your environment, you can take advantage of integration components provided by TIBCO BusinessEvents. Integration between TIBCO BusinessEvents and ActiveMatrix BusinessWorks enables each product to take advantage of the strengths of the other. For example, ActiveMatrix BusinessWorks can use TIBCO BusinessEvents as a light-weight rules engine, and TIBCO BusinessEvents can use transports available in ActiveMatrix BusinessWorks.

This guide assumes you know how to work with ActiveMatrix BusinessWorks processes. See *TIBCO BusinessWorks Process Design Guide* if you need detailed guidance.

Integration components include an ActiveMatrix BusinessWorks plug-in, which provides a palette of TIBCO BusinessEvents activities, and various TIBCO BusinessEvents functions. See [Integration Components on page 639](#) for details. You can use these components to do the following:

- More easily send and receive TIBCO BusinessEvents events using ActiveMatrix BusinessWorks activities.
- Invoke TIBCO BusinessEvents rule functions from ActiveMatrix BusinessWorks
- Call ActiveMatrix BusinessWorks processes from TIBCO BusinessEvents.

This chapter explains how to configure the environment and use these features.



See the product readme for the supported version of TIBCO ActiveMatrix BusinessWorks.

TIBCO BusinessEvents requires JRE 6. You must also use JRE 6 with ActiveMatrix BusinessWorks in order to integrate with TIBCO BusinessEvents.

Integration Components

Integration components include a palette of activities in ActiveMatrix BusinessWorks, and a category of functions in TIBCO BusinessEvents.

ActiveMatrix BusinessWorks Activities Palette

RuleServiceProvider Configuration Resource

The RuleServiceProvider Configuration resource is used to identify the source of the ontology definitions (in an EAR file) and the location of the TIBCO BusinessEvents application at runtime.

Event Related Activities

The Receive Event, Send Event, And Wait for Event activities enable you to work more easily with TIBCO BusinessEvents events within ActiveMatrix BusinessWorks. The ontology definitions from the TIBCO BusinessEvents application are available to these activities so you can easily configure them. See the following sections:

- [Receive Event Resource Reference on page 646](#)
- [Send Event Resource Reference on page 648](#)
- [Wait for Event Resource Reference on page 649](#)

Invoke RuleFunction Activity

An Invoke RuleFunction activity invokes a TIBCO BusinessEvents rule function in a specified agent instance and passes information to it. The Invoke RuleFunction activity can be combined with any process starter. Execution is synchronous.

TIBCO BusinessEvents Functions

TIBCO BusinessEvents provides a BusinessWorks category of functions, as follows:

BusinessWorks.invokeProcess(): Invokes an ActiveMatrix BusinessWorks process in synchronous mode and waits for completion of the process before returning to the rule or rule function. Starts the process engine if it is not already started. Returns an event, or null. Generates an advisory event if it times out.

BusinessWorks.startProcess(): Invokes an ActiveMatrix BusinessWorks process in asynchronous mode and returns the process ID (job ID). Upon completion, the ActiveMatrix BusinessWorks process passes an event to the rule function specified in an argument of `startProcess()`. Starts the process engine if it is not already started.

BusinessWorks.cancelProcess(): Cancels the specified ActiveMatrix BusinessWorks process. Useful for canceling a long running job. (Cancellation is not guaranteed because the process may complete before receiving the cancellation command.)

BusinessWorks.init(): Used as needed to initialize the ActiveMatrix BusinessWorks engine.

BusinessWorks.shutdown(): Used as needed to stop the ActiveMatrix BusinessWorks engine.

Design Considerations

This section presents some considerations to keep in mind when you are designing your integration project.

Integration Scope

You can only use TIBCO BusinessEvents resources that are included in the TIBCO BusinessEvents EAR file referenced in the RuleServiceProvider Configuration resource.

TIBCO BusinessEvents and ActiveMatrix BusinessWorks must be installed on the same machine.

Thread Management

TIBCO BusinessEvents generates its own threads to execute the rule function that the ActiveMatrix BusinessWorks process calls through the Invoke RuleFunction activity. Then the ActiveMatrix BusinessWorks thread is released and the process is set to a pending state. When the rule function returns, the ActiveMatrix BusinessWorks process resumes its Ready state.

Avoiding Threading Issues

It is possible to configure a complex execution path in your project, for example, one in which TIBCO BusinessEvents invokes an ActiveMatrix BusinessWorks process, which in turn invokes a TIBCO BusinessEvents rule function, and so on.

The `BusinessWorks.InvokeProcess()` rule function and the Invoke RuleFunction activity both operate synchronously. Both attempt to acquire a lock on the TIBCO BusinessEvents RTC.

Check your execution path carefully to ensure that there are no threading issues leading to deadlock. In general, make sure that no action in the entire execution path attempts to use the same working memory that is already locked.

For example, if you execute `BusinessWorks.InvokeProcess()` function in TIBCO BusinessEvents, then in the ActiveMatrix BusinessWorks process it calls, you cannot use an Invoke RuleFunction activity to invoke a rule function in the same agent. You could invoke a rule function running in a different agent, however.

Configuring a RuleServiceProvider Configuration Resource

You reference a RuleServiceProvider Configuration resource in various activities so ActiveMatrix BusinessWorks can locate the TIBCO BusinessEvents resources used in the integration.

1. Open the ActiveMatrix BusinessWorks project in TIBCO Designer and open the project folder where you want to store the resource.
2. Right-click in the design panel and select **Add Resource > TIBCO BusinessEvents Activities > TIBCO BusinessEvents RuleServiceProvider Configuration**.
3. Enter values in the Configuration tab, following guidelines in [TIBCO BusinessEvents RuleServiceProvider Configuration Resource Reference on page 643](#).
4. **Click** Apply and save the resource.



A RuleServiceProvider Configuration shared resource is also available In TIBCO BusinessEvents Studio. It can be used during migration to display a migrated version 3.x resource for informational purposes.

TIBCO BusinessEvents RuleServiceProvider Configuration

Resource Reference



This resource specifies which TIBCO BusinessEvents project to use for the integration. See also [Configuring a RuleServiceProvider Configuration Resource on page 642](#).

Configuration

The Configuration tab has the following fields.

Field	Global Var?	Description
Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
Description	No	Short description of the resource.
Repo URL/EAR Path	Yes	<p>One of the following:</p> <p>Path to and name of the TIBCO BusinessEvents EAR file. Ontology definitions in this EAR file are used by activities in the TIBCO BusinessEvents activities palette.</p> <p>(At design or testing time you can use an exact copy of the EAR that is used for production. You can use a global variable so that you can define the location of the EAR at deploy time.)</p> <p>or</p> <p>The server-based repository URL for a TIBCO BusinessEvents project that was deployed to a TIBCO Administrator domain.</p>

Field	Global Var?	Description
Deployment Descriptor Path	Yes	<p>Path to the CDD (Cluster Deployment Descriptor) file used by the deployed TIBCO BusinessEvents application.</p> <p>At design or testing time you can use an exact copy of the CDD that is used for production. You can use a global variable so that you can define the location of the CDD at deploy time.</p>
Processing Unit ID		ID of a processing unit in the specified CDD whose values are used for the integration.

Working With the TIBCO BusinessEvents Activities

This guide assumes you know how to work with ActiveMatrix BusinessWorks processes. See *TIBCO BusinessWorks Process Design Guide* if you need detailed guidance.



The TIBCO BusinessEvents Activities palette allows you to easily send and receive TIBCO BusinessEvents events in TIBCO ActiveMatrix BusinessWorks applications.

Before you work with other activities, configure a RuleServiceProvider Configuration resource so that the ActiveMatrix BusinessWorks project can locate TIBCO BusinessEvents resources. See [TIBCO BusinessEvents RuleServiceProvider Configuration Resource Reference on page 643](#).

To Work with the TIBCO BusinessEvents Activities

1. Open the folder in which you want to add a process definition. Right-click in the design panel, and select **Add Resource > Process > Process Definition**.
2. Double-click the Process Definition resource to open it. You see the Start and End activities.
3. Right-click in the design panel, and select **Add Resource > TIBCO BusinessEvents Activities**. Select activities as needed to configure the process. See the following for guidelines:
 - [Receive Event Resource Reference on page 646](#)
 - [Send Event Resource Reference on page 648](#)
 - [Wait for Event Resource Reference on page 649](#)
 - [Invoke RuleFunction Resource Reference on page 656](#)
4. Click **Apply** and save the resource.

Receive Event Resource Reference



This process starter starts a process when the specified simple event arrives from its default destination channel, or from a destination specified in the configuration tab.

Configuration

The Configuration tab has the following fields.

Field	Global Var?	Description
Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
Description	No	Short description of the resource.
RuleServiceProvider Configuration		Browse to and select the RuleServiceProvider Configuration resource used in this process. The lists of events and destinations shown in the Event Reference and Custom Destination fields comes from the EAR file referenced in the selected RuleServiceProvider Configuration resource.
Event Reference	No	Browse to and select the simple event you want to receive.
Custom Destination	No	If you do not want to listen for the specified event on its default destination, browse to and select an alternate destination.

Misc

The Misc tab has the following fields.

Field	Description
Sequencing Key	<p>This field can contain an XPath expression that specifies which processes should run in order. Process instances with sequencing keys that evaluate to the same value will be executed sequentially in the order the process instance was created.</p> <p>See <i>TIBCO ActiveMatrix BusinessWorks Process Design Guide</i> for more information about controlling the execution order of process instances and about XPath expressions.</p>
Custom ID	<p>This field can contain an XPath expression that specifies a custom ID for the process instance. This ID is displayed in the View Service dialog of TIBCO Administrator, and it is also available in the <code>\$_processContext</code> process variable.</p>

Output

The Output tab has the following fields.

Output Item	Data Type	Description
BEReceiveEvent Output	complex	Expand BEReceiveEventOutput to show the name, properties, and attributes of the event received.

Send Event Resource Reference



Data from the ActiveMatrix BusinessWorks process context is used to send an event of a specified type to its default destination or to a destination specified in the configuration tab.

Configuration

The Configuration tab has the following fields.

Field	Global Var?	Description
Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
Description	No	Short description of the resource.
RuleServiceProvider Configuration		Browse to and select the RuleServiceProvider Configuration resource used in this process. The lists of events and destinations shown in the Event Reference and Custom Destination fields comes from the EAR file referenced in the selected RuleServiceProvider Configuration resource.
Event Reference	No	Browse to and select the simple event you want to send.
Custom Destination	No	If you do not want to send the specified event to its default destination, browse to and select an alternate destination.

Input

The Input tab has the following fields.

Input Item	Data Type	Description
BESendEventInput	complex	Expand BESendEventInput to show the name, properties, and attributes of the event type selected in the Configuration tab.

Wait for Event Resource Reference



This activity listens to the default destination of a specified event type or to a destination specified in the configuration tab, and waits for a matched simple event.

Configuration

The Configuration tab has the following fields.

Field	Global Var?	Description
Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
RuleServiceProvider Configuration		Browse to and select the RuleServiceProvider Configuration resource used in this process. The lists of events and destinations shown in the Event Reference and Custom Destination fields comes from the EAR file referenced in the selected RuleServiceProvider Configuration resource.
Description	No	Short description of the resource.
Event Reference	No	Browse to and select the simple event you want to wait for.
Custom Destination	No	If you do not want to listen for the specified event on its default destination, browse to and select an alternate destination.

Event

The Event tab has the following fields.

Input Item	Global Var?	Description
Candidate Event Key	No	<p>Use to filter events. Only events whose candidate event key matches the key provided will trigger the activity.</p> <p>For example you can use the event's extID as the event key. You can also use any of the event's properties as a matching candidate event key. There can be multiple keys, but only the first one that matches is considered.</p> <p>As an example use of this property, suppose you want "Wait for Event" to fire only for the 10th occurrence of an event. You would set the Candidate Event Key value to 10. In the input tab, you would map the input item key to a global variable that is incremented by 1 each time the event occurs.</p> <p>Refer to <i>ActiveMatrix BusinessWorks Palette Reference</i> sections on other "Wait For" types of activities (for example, for JMS queues or topics, or for Rendezvous) for more information.</p>
Event Timeout	No	<p>If the event is received before the process reaches this activity, the event waits for this number of milliseconds. If the timeout period ends before the process reaches this activity, the event is discarded.</p>

Input

The Input tab has the following fields.

Input Item	Data Type	Description
key	String	Candidate Event Key value (specified on the Event tab
processTimeout	Integer	The processTimeout value is the Event Timeout from Event tab

Output

The Output tab has the following fields.

Input Item	Data Type	Description
BEReceiveEventOutput	complex	Expand BEReceiveEventOutput to show the name, properties, and attributes of the event type selected in the Configuration tab.

Invoking a TIBCO BusinessEvents Rule Function from a Process

To enable an ActiveMatrix BusinessWorks process to call a TIBCO BusinessEvents rule function, you need to configure an Invoke RuleFunction activity.

See [Integration Components on page 639](#) for an introduction to the integration features.



- See [Thread Management on page 641](#) for important information about avoiding deadlock.
- Be careful in using the `ThreadLocal` variable because the thread of execution may not be same as that in the prior activity.
- Concepts passed to the rule function are not automatically asserted. Similarly, events passed to the rule function are not automatically asserted. You must assert them explicitly.
- `Object` support is not provided. You can't set `Object` in the scope of the rule function or as its return value. You must specify a specific type of TIBCO BusinessEvents object or a primitive (other than `Object`).

Specifying Input Arguments

All TIBCO BusinessEvents primitives (except `Object`) are supported. Appropriate boxing and unboxing of XML primitive types to TIBCO BusinessEvents primitive types is handled. For more details, see [Chapter 20, Rule Language Datatypes, on page 339](#).

You can pass null values. Select the argument in the Activity Input panel of the Input tab. Click the Edit Statement (exclamation point) button. In the Edit Statement dialog, Content tab, check the Set Explicit Nil checkbox. If `xsi:nil==true`, then a null is passed as the argument.

Base Concept and Base Event are supported as argument types. You must map a specific type to the base concept or base event. Open the Edit Statement dialog, and in the Type tab, check the Type Substitution check box. In the Type field, select XML Type Reference. Select a Schema and a Type.

Using Synchronous Invocation

Execution is synchronous. The process waits for the rule function to return a value. The rule function acquires a lock on TIBCO BusinessEvents working memory until the RTC cycle completes and then returns a value to the process.

You can use more than one Invoke RuleFunction activity in a process only if each uses a different agent. See [Thread Management on page 641](#) for more details.

Using the lockWM Parameter

By default the working memory is locked during the invocation (that is, the value is true). You can change this behavior by setting the `lockWM` parameter to false.

When you set the value to false, the rule function executes outside the context of the working memory.

The results must not modify concept instances or scorecards (as in the case of a preprocessor).

Set the `lockWM` parameter to false for optimization purposes only and use with care. Set it to false only if the call does not modify the working memory in any way. For example, use it to perform stateless calls or calculations, or to create new concept instances or events, or for lookups and so on.

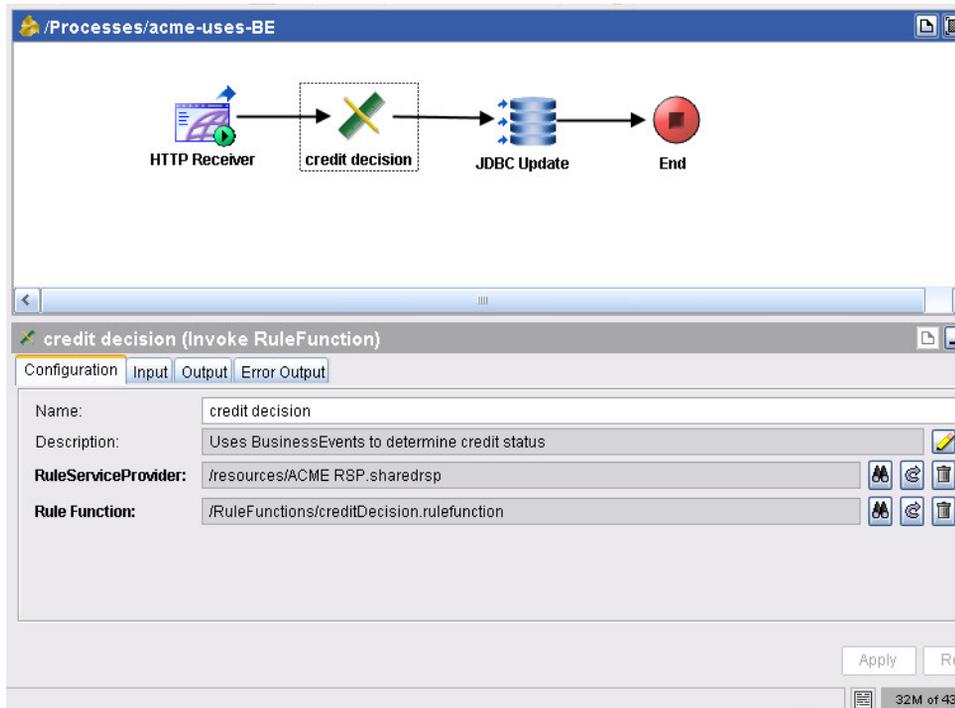
Overriding the Rule Function at Runtime

You can override which function is called at runtime. To do so, enter a global variable as the value for `rulefunction` in the input tab. At runtime, you can specify a rule function with a different name but the same signature as the one specified in the activity's Configuration tab.

You can only use rule functions that are in the TIBCO BusinessEvents EAR file deployed for the integration project.

Working With Invoke RuleFunction Activities

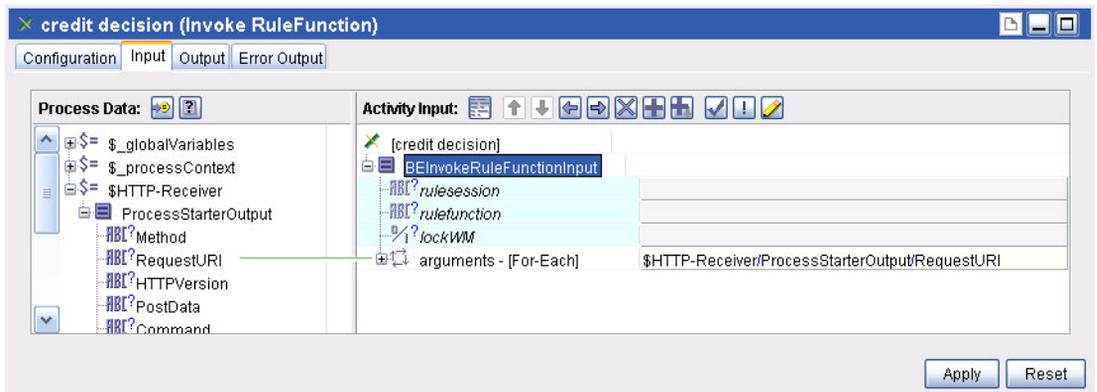
See [Invoking a TIBCO BusinessEvents Rule Function from a Process on page 652](#) for important information about using the Invoke RuleFunction activity in your ActiveMatrix BusinessWorks process.



To Configure an Invoke RuleFunction Activity

1. Start TIBCO Designer and open (or add) the ActiveMatrix BusinessWorks process in which you want to use an Invoke RuleFunction activity.
2. Right-click in the design panel and select **Add Resource > TIBCO BusinessEvents Activities > Invoke RuleFunction**. Link the activity to other activities as appropriate.

3. Complete the Configuration tab, following guidelines provided in [Invoke RuleFunction Resource Reference](#) on page 656.
 - Name the activity and give it a description as desired.
 - Browse to and select the RuleServiceProvider Configuration resource that contains the ontology for the project you are integrating with.
 - Browse to and select the TIBCO BusinessEvents rule function you want to invoke.
4. Select the **Input** tab. Expand each item to see all fields. Enter values or map values from process data. Follow guidelines provided in [Invoke RuleFunction Resource Reference](#) on page 656.



5. Click **Apply** and save the resource.

Invoke RuleFunction Resource Reference



To enable an ActiveMatrix BusinessWorks process to call a TIBCO BusinessEvents rule function, you configure one or more Invoke RuleFunction activities according to your needs. Execution is synchronous. See [Invoking a TIBCO BusinessEvents Rule Function from a Process on page 652](#) and [Working With Invoke RuleFunction Activities on page 654](#) for more details.

Configuration

The Configuration tab has the following fields.

Field	Global Var?	Description
Name	No	The name to appear as the label for the resource. Names follow Java variable naming restrictions. Do not use any reserved words. Names must be unique within a folder. See Identifier Naming Requirements on page 323 .
Description	No	Short description of the resource.
RuleServiceProvider Configuration		Browse to and select the RuleServiceProvider Configuration resource used in this process. The list rule functions you can select for the Rule Function field comes from the EAR file referenced in the selected RuleServiceProvider Configuration resource.
Rule Function	No	Select the desired rule function from the TIBCO BusinessEvents project.

Input

The Input tab has the following fields.

Field	Global Var?	Description
rulesession	Yes	Specifies the agent where the rule function is to execute. Optional if the processing unit has only one agent To specify a value, enter the agent name (as specified in the CDD file).
rulefunction	Yes	Optional. Allows you to override the rule function specified in the Configuration tab. Enter a global variable that is used to specify the name of the rule function at runtime. Do not enter the .rulefunction extension. The rule function you specify must have the same signature as the function specified in the Configuration tab (and a different name). See Overriding the Rule Function at Runtime on page 653
lockWM	Yes	Default value is true. The working memory is locked during the invocation. You can set the value to false only if conditions explained in the section Using the lockWM Parameter on page 653 are met.
arguments	No	Displays the arguments for the rule function, if it has any. Map the available process data to the activity input, or enter values in the fields as appropriate for the datatype of each argument.

Output

The Output tab has the following fields.

Output Item	Data Type	Description
Return	(Varies)	The return type for the specified rule function. Return can be void or any of the following types: String, integer, long, double, Boolean, datetime, concept, event.

Working with the BusinessWorks Functions

The integration model provides a BusinessWorks category of functions. This category is located in the General Functions library. Use and configuration of the functions in the BusinessWorks category are described in this section. See [Integration Components on page 639](#) for a brief introduction to these features.

Providing Paths to TIBCO BusinessEvents Project Resources Using Schemas

The `BusinessWorks.invokeProcess()` and `BusinessWorks.startProcess()` functions can send events to and receive events from ActiveMatrix BusinessWorks processes. However, the paths to TIBCO BusinessEvents project resources are not available by default in TIBCO Designer. To use these functions successfully, you must generate the event schemas in TIBCO BusinessEvents and import them into the TIBCO Designer project. See [Exporting \(Generating\) Concept and Event Schema \(XSD\) Files on page 5](#) for details on exporting the event schemas.

In the TIBCO Designer project, import these schemas and use them in the relevant processes, so that the event resource path is known to TIBCO Designer. See ActiveMatrix BusinessWorks documentation for details.

Using `invokeProcess()`

Purpose Use this function to take advantage of ActiveMatrix BusinessWorks features. For example, you could use `invokeProcess()` to send customer information to an ActiveMatrix BusinessWorks process that gets the discount information from a database and returns it to TIBCO BusinessEvents.

Runtime Behavior If the ActiveMatrix BusinessWorks engine is not already started, TIBCO BusinessEvents will start it before `invokeProcess()` executes. See [Using `init\(\)` on page 662](#) for another way to start the engine.

When a rule containing or calling `invokeProcess()` is triggered, `invokeProcess()` calls the specified ActiveMatrix BusinessWorks process and passes it an event. The `invokeProcess()` function executes synchronously, so the TIBCO BusinessEvents rule engine waits. The ActiveMatrix BusinessWorks process performs its work and returns an event or null at completion of the process, or it times out if you set a timeout that is exceeded. It generates an advisory event if it times out. Events received are not automatically asserted.

Error Handling and Advisory Events

TIBCO BusinessEvents asserts an advisory event and returns null if the ActiveMatrix BusinessWorks process fails, or if the invocation times out. The advisory event category is `Engine`, and the type is `INVOKE_BW_PROCESS`. The message contains the error message from the failed ActiveMatrix BusinessWorks process, or the timeout message. The advisory event is also created if the process is cancelled using `cancelProcess()`, because the system cannot differentiate between different causes for the process stopping before completion.

Because the `invokeProcess()` function returns null when an error occurs, you must handle the possibility of a null return in rules that use any property or attribute of the returned event.

See [Chapter 11, Advisory Events, on page 187](#) for more information about advisory events.



Do not use `invokeProcess()` more than once in the same thread of execution. See [Thread Management on page 641](#).

Configuring the Function

When configuring the `invokeProcess()` function in a rule or rule function, specify the parameters as follows:

- The ActiveMatrix BusinessWorks process that the function invokes. The specified process must not contain a process starter.
- The event to pass to the process (or specify null if you don't want to pass an event to the process).
- A timeout (or specify zero if you don't want to use a timeout).

For example, you send customer information from an event (alias `neworder`) to an ActiveMatrix BusinessWorks process, which returns an event with the discount level to offer:

```
Events/Discountlevel Discount;
Discount = Businessworks.invokeProcess("/Processes/CustInfo",
neworder, 0);
```



Timeouts If you set a timeout period and a timeout occurs, the rule or rule function containing `InvokeProcess()` continues without waiting for the ActiveMatrix BusinessWorks process to complete. If you use a timeout, set it to a period long enough for the ActiveMatrix BusinessWorks process to complete. Include logic to handle the case that the `invokeProcess()` function does time out.

Configuring the Process

Configure the Start activity of the specified ActiveMatrix BusinessWorks process to accept the event passed by the `invokeProcess()` function, if an event is passed. If you specify null, of course, this step is not required.

Configure the rest of the activities in the process to carry out whatever processing is desired using the data passed into it by `invokeProcess()`.

Similarly, in the End activity, specify the event type to return to the `invokeProcess()` process.

The returned event is then used as needed by the logic of the rule or rule function.



Events returned are not asserted. You must explicitly assert them as needed.

See [Providing Paths to TIBCO BusinessEvents Project Resources Using Schemas](#) above for additional setup you must do.

Using `startProcess()`

Purpose Use the `startProcess()` function when you want to invoke an ActiveMatrix BusinessWorks process that performs work that can be completed asynchronously. The `startProcess()` function invokes an ActiveMatrix BusinessWorks process in asynchronous mode and immediately returns the job ID of the process. Rule processing continues. When the ActiveMatrix BusinessWorks process completes, it passes an event to a callback rule function that is specified in a `startProcess()` argument. For example, you send order information to an ActiveMatrix BusinessWorks order fulfillment process. When the order ships, notification is returned to TIBCO BusinessEvents, which updates a customer concept instance.

Runtime Behavior If the ActiveMatrix BusinessWorks engine is not already started, TIBCO BusinessEvents will start it before `startProcess()` executes. See [Using `init\(\)` on page 662](#) for another way to start the engine.

When a rule containing or calling `startProcess()` is triggered, `startProcess()` calls the specified ActiveMatrix BusinessWorks process and passes it an event. `startProcess()` returns the `jobID` of the ActiveMatrix BusinessWorks process. (You can use the returned job ID, for example, in the `cancelProcess()` rule function.) The function executes asynchronously, so the TIBCO BusinessEvents rule engine continues while at the same time the ActiveMatrix BusinessWorks process executes. The ActiveMatrix BusinessWorks process performs its work and passes an event to the callback (`ruleFnURI`) rule function specified in the `startProcess()` function arguments. The `ruleFnURI` rule function performs its work, for example, creating and asserting the event.



The `ruleFnURI` rule function must not modify concept instances or scorecards.

- Configuring the Function** When configuring the `startProcess()` function in a rule or rule function, specify the following in the parameters:
- The ActiveMatrix BusinessWorks process that the function invokes. The specified process must not contain a process starter.
 - The event to pass to the process (or specify null if you don't want to pass an event to the process).
 - The TIBCO BusinessEvents rule function that the process calls on completion (the callback rule function). The required signature for the callback rule function is shown below.
- Configuring the ruleFnURI Rule Function** The callback rule function, specified in the `startProcess()` `ruleFnURI` argument, is called when the ActiveMatrix BusinessWorks process completes. Add the `ruleFnURI` rule function to the TIBCO BusinessEvents project.
- For convenience, this rule function is referred to as the `ruleFnURI` rule function in this section. The `ruleFnURI` rule function must have the following signature, and the Validity field (in the Configuration tab) must be set to Action:
- ```
void ruleFn(long jobID, int status, Event outputEvent, Object closure)
```
- **jobID** Type: long. The job id can be used to correlate the information passed to `ruleFnURI` with related information in TIBCO BusinessEvents.
  - **status** Type: int. Returns 0 (zero) if the process completed successfully, and -1 if the process did not complete successfully (for example because `cancelProcess()` was called).
  - **outputEvent** Type: Event. An event passed to the rule function by the ActiveMatrix BusinessWorks process. It can be created by the process, or it can be an existing event. If the process fails or is cancelled (for example because `cancelProcess()` was called), an advisory event is returned.
  - **closure** Type: Object. A closure object could be, for example, a value from the original context that has to be passed back to the TIBCO BusinessEvents engine, for example, a loan rate that has been promised. (Note that type is Object, so you cannot pass an event or concept.)
- Configuring the Process** Configure the Start activity of the specified ActiveMatrix BusinessWorks process to accept the event passed by the `startProcess()` function, if an event is passed. If you specify null, of course, this step is not required.
- Configure the rest of the activities in the process to carry out whatever processing is desired using the data passed into it by `startProcess()`.
- In the End activity, specify the event type that is to be passed to the specified `ruleFnURI` rule function.

See [Providing Paths to TIBCO BusinessEvents Project Resources Using Schemas on page 658](#) above for additional setup you must do.

The returned event is then used as needed by the logic of the rule or rule function. Note that if you want to assert the event that is returned, you must explicitly assert it.

## Using cancelProcess()

**Purpose** Use the `cancelProcess()` function to cancel a long running process, specified by `jobID`.

Cancellation may fail if the process has already completed before receiving the cancellation command. In this case, the following exception is thrown:

```
java.lang.Exception: Job JobId not found
```

**Configuration** The job ID for the process you want to cancel is provided in the return value of `startProcess()`.

## Using init()

**Purpose** The `init()` function initializes the ActiveMatrix BusinessWorks engine if it is not already running. Use of `init()` when the engine is already running is harmless.

Use of `init()` is optional. It is provided as a convenience. For example you can use it in a TIBCO BusinessEvents startup rule function to initialize the ActiveMatrix BusinessWorks engine at startup, so as to save valuable time later.

If `invokeProcess()` or `startProcess()` are executed when the ActiveMatrix BusinessWorks engine is not already started, they will start the engine at that time.

**Configuration** A general good practice is to call `init()` to start the ActiveMatrix BusinessWorks engine in a startup rule function (specified in the Startup/Shutdown tab of the BAR). The `init()` function takes no parameters. The rule function might simply contain this code:

```
BusinessWorks.init();
```

## Using shutdown()

**Purpose** The `shutdown()` function shuts down the ActiveMatrix BusinessWorks engine if it is running. Use of `shutdown()` when the engine is already shut down is harmless.

Use of `shutdown()` is optional. It is recommended that you use it only when you have finished using the ActiveMatrix BusinessWorks engine, and won't need to start it again. Stopping and restarting the engine is not necessary and can affect performance

**Configuration** Simply place the shutdown command so that it is executed when needed. The `shutdown()` function takes no parameters. The rule function might simply contain this code:

```
BusinessWorks.shutdown();
```



# TIBCO BusinessEvents Performance Profiler

This chapter explains how you can run a profiler utility to gather statistics about activities that occur during each RTC cycle. This information helps to identify bottlenecks in the project, which can often be addressed by redesigning rules or other aspects of a project.

## Topics

---

- [Overview of Profiler, page 666](#)
- [Changing the Delimiter Character, page 667](#)
- [Turning Profiler On and Off, page 668](#)
- [Profiler Reference, page 674](#)

## Overview of Profiler

---

The profiler utility collects statistics relating to the run to completion (RTC) rule evaluation cycle in an inference agent. The utility does not collect data about object management. It also does not collect data for other types of agents. It helps you understand, for example, internal execution times and frequencies of rules.

The profiler records time spent during each RTC on activities such as number of times each condition or action is performed, and total time spent on each condition and action. A complete RTC includes conditions and actions, although any individual RTC might contain only conditions or only actions.

Statistics are collected for each completed RTC. When the profiler is directed to stop during an RTC, it continues to collect data for the current RTC until that RTC is completed.

When the profiler is turned off, it continues to write statistics for the current session until that session is completed. So the RTC in progress is always completed, even if the profiler is directed to stop during an RTC.

After the profiler finishes, statistics data is written to the specified (or default) file and cleared from memory. The engine continues to run. If the engine stops before the profiler completes, the file is not created.

You can execute the profiler and turn it off in three ways:

**Using properties** Profiler turns on when the agent initializes at system startup. Used to profile RTC time, including startup rule functions.

**Using TIBCO BusinessEvents catalog functions** Profiler turns at the beginning of the next RTC after the function call, if it has not already been enabled. (There is no effect if the profiler is already on). Used to turn the profiler on and off inside a rule or rule function.)

**Using a TIBCO Hawk method** Profiler turns on by invoking a TIBCO BusinessEvents microagent Hawk method. Profiler is turned on at the beginning of the next RTC after the method call, if it has not already been enabled. (There is no effect if the profiler is already on). Used to dynamically turn the profiler on and off.

**Using a TIBCO BusinessEvents Monitoring and Management method** Similar to using a Hawk method.

See [Turning Profiler On and Off on page 668](#) for details.

See [Profiler Reference on page 674](#) for a reference to all output data file column headings.

## Changing the Delimiter Character

---

The profiler is tab-delimited by default. The delimiter character can be changed adding the following property in the CDD file:

```
be.engine.profile.delimiter
```

Specify the delimiter using a String value. Enclose the value in double quotes (the quotes are not used as part of the delimiter).

For example to use an open curly brace as the delimiter, you would specify "{" as the value. Do not choose a character used in rule conditions.

Use a single character if the application into which you will import the output uses a one-character delimiter. When importing the file into Excel, do not check the "Treat consecutive delimiters as one" option. Consecutive delimiters indicate a column that is empty.



Also, when importing the file into Excel, set the timestamp field to Text (and not General, which is the default).

## Turning Profiler On and Off

---

This section explains the different ways you can turn the profiler on and off.

### To Turn Profiler On and Off Using TIBCO BusinessEvents Monitoring and Management

If you have deployed the processing unit using TIBCO BusinessEvents Monitoring and Management (MM), you can turn the profiler on and off using the MM Console.

Use of MM is documented in the *TIBCO BusinessEvents Administration* guide. In particular, see the section *To Execute a Method* and the section *Profiler Group*, in Chapter 6, *Monitoring and Managing a TIBCO BusinessEvents Cluster with MM*.

### To Turn Profiler On and Off Using Properties

Set the following properties in the Cluster Deployment Descriptor (CDD) Processing Unit tab, for all processing units (engines) whose RTC performance you want to profile.

Table 66 Profiler Configuration Properties (Sheet 1 of 3)

| Property                                               | Notes                                                                                                                                                                    |
|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>be.engine.profile.Agent_Class_Name.enable</code> | <p>If set to true, enables profiler for the specified agent class (<i>Agent_Class_Name</i>) when the agent initializes.</p> <p>Default is false.</p>                     |
| <code>be.engine.profile.*.enable</code>                | <p>If set to true, enables the profiler for all agents when each agent initializes, even when a specified agent class profiler is disabled.</p> <p>Default is false.</p> |

Table 66 Profiler Configuration Properties (Sheet 2 of 3)

| Property                                                 | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>be.engine.profile.Agent_Class_Name.file</code>     | <p>The name of output file that the profiler writes to, for the specified agent (<i>Agent_Class_Name</i>).</p> <p>Default behavior is as follows:</p> <p>If <code>be.engine.profile.*.file</code> is specified and <code>be.engine.profile.Agent_Class_Name.file</code> is not specified, then the file name is the value of <code>be.engine.profile.*.file</code>, with the <i>Agent_Class_Name</i> appended.</p> <p>If the properties <code>be.engine.profile.*.file</code> and <code>be.engine.profile.Agent_Class_Name.files</code> are not specified the name is created as follows: <code>be-profile_Agent_Class_Name.csv</code></p> |
| <code>be.engine.profile.*.file</code>                    | <p>The default (prefix for the) name of the output file that the profiler writes to. In all cases, the appropriate <i>Agent_Class_Name</i> is appended.</p> <p>Default name is <code>be-profile.csv</code> and it is located under the current working directory, if file name is not specified.</p>                                                                                                                                                                                                                                                                                                                                       |
| <code>be.engine.profile.Agent_Class_Name.duration</code> | <p>Specifies the duration of profile data collection in seconds, for the specified <i>Agent_Class_Name</i>).</p> <p>When the duration period ends, the profiler continues to collect statistics for the current RTC until the RTC is completed, then outputs data and stops. So the RTC in progress is always completed, even if the profiler is directed to stop during an RTC.</p> <p>If you set duration to a value of zero or less (<math>\leq 0</math>), then profiling continues until agent stops or profiler is explicitly turned of using a function or Hawk method.</p> <p>Default is -1.</p>                                    |

Table 66 Profiler Configuration Properties (Sheet 3 of 3)

| Property                                              | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>be.engine.profile.*.duration</code>             | <p>Specifies the duration of profile data collection in seconds, for all agents.</p> <p>When the duration period ends, the profiler continues to collect statistics for the current RTC until the RTC is completed, then outputs data and stops. So the RTC in progress is always completed, even if the profiler is directed to stop during an RTC.</p> <p>If you set duration to a value of zero or less (<math>\leq 0</math>), then profiling continues until the agents stop or profiler is explicitly turned of using a function or Hawk method.</p> <p>When <code>be.engine.profile.Agent_Class_Name.duration</code> and <code>be.engine.profile.*.duration</code> are both present, the duration specified in <code>be.engine.profile.Agent_Class_Name.duration</code> takes precedence.</p> <p>Default is -1.</p> |
| <code>be.engine.profile.Agent_Class_Name.level</code> | <p>Level of depth that profile data will be collected for the specified agent (<i>Agent_Class_Name</i>):</p> <ul style="list-style-type: none"> <li>• -1: all levels of profile data are collected, including RTC level and conditions and actions within the RTC.</li> <li>• 1: Only RTC level of profile data will be collected (and no condition and action data).</li> </ul> <p>Default is -1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <code>be.engine.profile.*.level</code>                | <p>Level of depth that profile data will be collected for all agents:</p> <ul style="list-style-type: none"> <li>• -1: all levels of profile data are collected, including RTC level and conditions and actions within the RTC.</li> <li>• 1: Only RTC level of profile data will be collected (and no condition and action data).</li> </ul> <p>When <code>be.engine.profile.Agent_Class_Name.level</code> and <code>be.engine.profile.*.level</code> are both present, the level specified in <code>be.engine.profile.Agent_Class_Name.level</code> takes precedence.</p> <p>Default is -1.</p>                                                                                                                                                                                                                         |

## To Turn Profiler On and Off Using Functions

This section assumes you understand how to use TIBCO BusinessEvents functions. It tells you which functions to use and the effect of each function.



You can turn the profiler on using the engine properties, and turn it off using a function or Hawk method, as desired. See notes for `be.engine.Agent_Class_Name.profile.duration` and `be.engine.*.profile.duration` in [Table 66, Profiler Configuration Properties, on page 668](#).

**To turn the profiler on** In your rule or rule function, use the following function to turn on the profiler:

---

```
Engine.Profiler.startCollectingToFile(String fileName, int level, long duration)
```

---

The above function turns on the TIBCO BusinessEvents Profiler and starts collecting data for the specified duration for the agent in which the rule or rule function that calls this function is executed. The profiler starts collecting data at the beginning of next RTC.

Profile data is output to the specified file in comma-separated value format at the end of the duration period, unless the profiler is turned off before the end of the duration, in which case it is output at the end of the RTC that completes after the profiler is turned off.

Input arguments are the same as the engine properties show in [Table 66, Profiler Configuration Properties, on page 668](#):

`String fileName`: See `be.engine.profile.Agent_Class_Name.file`

`int level`: See `be.engine.profile.Agent_Class_Name.level`

`long duration`: See `be.engine.profile.Agent_Class_Name.duration` (Here `Agent_Class_Name` is the current agent in which the rule or rule function that calls this function is executed.)

**To turn the profiler off** In your rule or rule function, use the following function to turn off the profiler:

---

```
Engine.Profiler.stopCollecting()
```

---

The above function turns off the TIBCO BusinessEvents profiler and writes the profile data to a file for the agent in which the rule or rule function that calls this function is included (the file is output at the end of the RTC that completes after the profiler is turned off). There is no effect if the profiler is not on.

## To Turn Profiler On and Off Using TIBCO Hawk Methods

This section assumes you understand how to use TIBCO Hawk methods. It tells you which methods to use and the effect of each method



You can turn the profiler on using properties, and turn it off using a function or Hawk method, as desired. See notes for `be.engine.profile.duration` in [Table 66, Profiler Configuration Properties, on page 668](#).

**Before you Begin** Ensure that the property `hawk.enabled` is set to true in the CDD at the cluster level before the TIBCO BusinessEvents engine starts.

**To turn the profiler on** Use the following method to turn on the profiler:

---

```
StartFileBasedProfiler(String session, String fileName, int level, long duration)
```

---

The above method turns on the TIBCO BusinessEvents profiler for the specified agent. The profiler starts collecting data at the beginning of next RTC for the specified duration.

This method works the same way as the `Engine.Profiler.startCollectingToFile()` function (see [To Turn Profiler On and Off Using Functions on page 671](#)), except that it requires you to specify an agent class.

Input arguments are the same as the engine properties shown in [Table 66, Profiler Configuration Properties, on page 668](#):

`String session`: If you want to monitor multiple agents, execute the method once for each, specifying the agent class name in each case. If there is only one agent, the session parameter is optional.

`String fileName`: See `be.engine.profile.Agent_Class_Name.file`

`int level`: See `be.engine.profile.Agent_Class_Name.level`

`long duration`: See `be.engine.profile.Agent_Class_Name.duration`

If you attempt to turn on the profiler when it is already running, an error is returned, but the running profiler is not affected.

**To turn the profiler off** In your rule or rule function, use the following function to turn off the profiler:

---

```
StopFileBasedProfiler(String session)
```

---

The above method turns off the TIBCO BusinessEvents profiler and writes the profile data into a file for the specified agent when the current RTC has completed. You must execute the method once for each session, as needed.

If you attempt to turn off the profiler when it is already off, an error is returned, but there is no effect on the profiler.

## Profiler Reference

The table in this section explains each of the columns in the profiler report. Data is grouped by `RTC_Stats_Type` and `Description`. (`Description` contains information about the specific RTC.) All data collected for conditions and actions performed during each RTC is listed within each RTC grouping.

Three rows of column headers for the RTC, condition and action are listed at the beginning of the file:

- One for statistics relating to the overall RTC
- One for statistics relating to conditions
- One for statistics relating to actions.

The first column of each data line is always the statistic type, which begins with one of `RTC-`, `CONDITION-`, or `ACTION-`.

Data is also grouped, one row for the overall RTC, and zero or more rows for different conditions or actions or both, as appropriate.

Table 67 Profiler Column Heading Reference (Sheet 1 of 5)

| Column Heading                                | Notes                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Statistics Relating to the Overall RTC</b> |                                                                                                                                                                                                                                                                                                          |
| <code>RTC_Stats_Type</code>                   | Type of rule evaluation cycle (RTC) There are 10 different types:<br>RTC:<br>RTC-Object-Asserted<br>RTC-Object-Modified<br>RTC-Object-Deleted<br>RTC-Event-Expired<br>RTC-Execute-Rule<br>RTC-Invoke-Action<br>RTC-Invoke-Function<br>RTC-Post-Process<br>RTC-Repeat-TimeEvent<br>RTC-Reevaluate-Element |
| <code>Timestamp</code>                        | The time at which the first RTC begins.                                                                                                                                                                                                                                                                  |
| <code>Description</code>                      | Information relating to the current <code>RTC_Stats_Type</code> . For example, the description of type <code>RTC-Object-Asserted</code> is the name of the object being asserted.                                                                                                                        |
| <code>NumExecuted</code>                      | Total number of times the same RTC has been executed.                                                                                                                                                                                                                                                    |

Table 67 Profiler Column Heading Reference (Sheet 2 of 5)

| Column Heading                    | Notes                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TotalRtcTime                      | Total time in milliseconds spent on the total number of executions of the same RTC.                                                                                                                                                                                                                                                     |
| AvgRtcTime                        | TotalRtcTime / NumExecuted                                                                                                                                                                                                                                                                                                              |
| MaxRtcTime                        | The maximum time in milliseconds spent on a single RTC.                                                                                                                                                                                                                                                                                 |
| MinRtcTime                        | The minimum time in milliseconds spent on a single RTC.                                                                                                                                                                                                                                                                                 |
| MaxResolvedTime                   | The maximum time in milliseconds spent to resolve a single RTC, including condition evaluation and action execution, but excluding operations related to object management (OM).                                                                                                                                                        |
| MinResolvedTime                   | The minimum time in milliseconds spent to resolve a single RTC, including condition evaluation and action execution, but excluding operations related to object management (OM).                                                                                                                                                        |
| Statistics Relating to Conditions |                                                                                                                                                                                                                                                                                                                                         |
| CONDITION_Stats_Type              | Type of rule Condition. One of the following:<br>CONDITION-Filter<br>Condition-Join                                                                                                                                                                                                                                                     |
| Timestamp                         | The timestamp of the first time the RTC begins.                                                                                                                                                                                                                                                                                         |
| RuleDescription                   | Name of the rule containing the condition, or name of state machine transition rule containing the condition.                                                                                                                                                                                                                           |
| ConditionDescription              | Condition statement of a rule or a state machine transition rule for user-defined condition, or predefined condition name for internal conditions.<br><br>When a user-defined rule condition has a commented-out line, the ConditionDescription of the next condition is<br><br><code>//... Only applies to CONDITION_Stats_Type</code> |
| NumEvaluated                      | Total number of times this condition is evaluated in the same RTC.                                                                                                                                                                                                                                                                      |
| NumEvalTrue                       | Total number of times the Join condition is evaluated to true.<br><br>This value is the sum of NumEvalTruePropagatedLeft and NumEvalTruePropagatedRight.                                                                                                                                                                                |

Table 67 Profiler Column Heading Reference (Sheet 3 of 5)

| Column Heading                 | Notes                                                                                                                                                        |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TotalTime                      | Total time in milliseconds spent on the total number of condition evaluations.                                                                               |
| AvgTime                        | $\text{TotalTime} / (\text{NumLeftSearch} + \text{NumRightSearch})$                                                                                          |
| MaxTime                        | The maximum time in milliseconds spent on a single condition evaluation.                                                                                     |
| MinTime                        | The minimum time in milliseconds spent on a single condition evaluation.                                                                                     |
| NumEvalPropagatedLeft          | Number of times the join condition evaluation is triggered by object assertion propagated from the left side of the condition.                               |
| NumEvalTruePropagatedLeft      | Number of times the join condition evaluates to true and evaluation is triggered by object assertion propagated from the left side of the condition.         |
| AvgRateEvalTruePropagatedLeft  | Average rate that the condition evaluates to true and evaluation is triggered by object assertion propagated from the left side of the condition.            |
| MaxNumEvalTruePropagatedLeft   | Maximum number of times the join condition evaluates to true and evaluation is triggered by object assertion propagated from the left side of the condition. |
| MinNumEvalTruePropagatedLeft   | Minimum number of times the join condition evaluate to true and evaluation is triggered by object assertion propagated from the left side of the condition.  |
| NumEvalPropagatedRight         | Number of times the join condition evaluation is triggered by object assertion propagated from the right side of the condition.                              |
| NumEvalTruePropagatedRight     | Number of times the join condition evaluates to true and evaluation is triggered by object assertion propagated from the right side of the condition.        |
| AvgRateEvalTruePropagatedRight | Average rate that the condition evaluates to true and evaluation is triggered by object assertion propagated from the right side of the condition.           |

Table 67 Profiler Column Heading Reference (Sheet 4 of 5)

| Column Heading                 | Notes                                                                                                                                                                                               |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MaxNumEvalTruePropagatedRight  | Maximum number of times the join condition evaluates to true and evaluation is triggered by object assertion propagated from the right side of the condition.                                       |
| MinNumEvalTruePropagatedRight  | Minimum number of times the join condition evaluate to true and evaluation is triggered by object assertion propagated from the right side of the condition.                                        |
| Statistics Relating to Actions |                                                                                                                                                                                                     |
| ACTION_Stats_Type              | Type of Actions.<br><br>There are 10 RTC types and four action types. The four action types are:<br><br>ACTION-Rule-Action<br>ACTION-Event-Expiry<br>ACTION-Invoke-Action<br>ACTION-Invoke-Function |
| Timestamp                      | The timestamp of first time the action execution begins.                                                                                                                                            |
| Description                    | Information about the action corresponding to current action type.<br><br>For example, description of type ACTION-Rule-Action is the name of the rule.                                              |
| NumExecuted                    | Total number of times the same action has been executed.<br><br>A complete action has two phases, action execution and operation.                                                                   |
| TotalActionTime                | Total time in milliseconds spent on the total number of actions.<br><br>$TotalActionTime = TotalExecutionTime + TotalOperationTime$                                                                 |
| AvgActionTime                  | $TotalActionTime / NumExecuted$ .                                                                                                                                                                   |
| MaxActionTime                  | The maximum time in milliseconds spent on a single action.                                                                                                                                          |
| MinActionTime                  | The minimum time in milliseconds spent on a single action.                                                                                                                                          |

Table 67 Profiler Column Heading Reference (Sheet 5 of 5)

| Column Heading     | Notes                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TotalExecutionTime | Total time in milliseconds spent on the total number of action execution phases. Action execution time is the time Rete network spends on executing the action, for example, time spent in creating new objects, deleting existing objects, and so on.                          |
| AvgExecutionTime   | $\text{TotalExecutionTime} / \text{NumExecuted}$ .                                                                                                                                                                                                                              |
| MaxExecutionTime   | The maximum time in milliseconds spent on a single action execution.                                                                                                                                                                                                            |
| MinExecutionTime   | The minimum time in milliseconds spent on a single action execution.                                                                                                                                                                                                            |
| TotalOperationTime | Total time in milliseconds spent on the total number of action operation phases. Action operation time is the time TIBCO BusinessEvents spends on applying changes to the Rete network, for example, time spent asserting newly created objects, or retracting deleted objects. |
| AvgOperationTime   | $\text{TotalOperationTime} / \text{NumExecuted}$ .                                                                                                                                                                                                                              |
| MaxOperationTime   | The maximum time in milliseconds spent on a single action operation.                                                                                                                                                                                                            |
| MinOperationTime   | The minimum time in milliseconds spent on a single action operation.                                                                                                                                                                                                            |
| MaxAgenda          | The maximum size of rule agenda as a result of all action operations.                                                                                                                                                                                                           |
| MinAgenda          | The minimum size of rule agenda as a result of all action operations.                                                                                                                                                                                                           |

This chapter explains how to test and debug projects within TIBCO BusinessEvents Studio.

The sections on debugger assume some familiarity with Eclipse Java debugger, as well as TIBCO BusinessEvents.



**Viewing the TIBCO BusinessEvents Studio Error Log** Errors in TIBCO BusinessEvents Studio functionality are reported in the error log. You can view the error log file in either of these ways:

- In TIBCO BusinessEvents Studio, navigate to Help > About TIBCO BusinessEvents Studio > Installation Details > Configuration > View Error Log
- In the file system, navigate to *Your\_Workspace* > *.metadata* > *.log*

## Topics

---

- [Overview of Testing and Debugging Projects, page 680](#)
- [Preparing to Run \(Test\) or Debug a Project, page 683](#)
- [Adding and Working with Launch \(Debug or Run\) Configurations, page 685](#)
- [Launch Configurations Reference, page 688](#)
- [Creating and Working With Test Data, page 690](#)
- [Setting Breakpoints in Rules and Rule Functions, page 693](#)
- [Running Debugger, page 695](#)
- [Running Tester, page 696](#)
- [Asserting Rule Input Data, page 697](#)
- [Viewing the Results, page 699](#)

## Overview of Testing and Debugging Projects

---

You can run and debug projects in TIBCO BusinessEvents Studio, using test data to understand how TIBCO BusinessEvents rules behave in response to inputs. Debugger shows potential problems in execution of rules and rule functions. You can also simply test a project by running it against test data. You can run multiple engines at the same time. For example, one engine's processing unit might run an inference agent while another's runs a cache agent.



The TIBCO BusinessEvents Decision Manager add-on product also uses a tester feature, Table Analyzer, to enable technical and non-technical users to test rules, rule functions, and decision tables. See TIBCO BusinessEvents Decision Manager documentation for details.

### Debugging Projects

TIBCO BusinessEvents Debugger allows you to control the execution of a project by setting breakpoints, stepping through your code, suspending launched programs, examining the contents of variables, providing rule input, and so on.

TIBCO BusinessEvents Debugger integrates with the Eclipse Java development toolkit debugger. Much of the functionality is standard Eclipse debugger functionality. See the Eclipse help for details on features such as breakpoint preferences and other functionality. This chapter explains only the TIBCO BusinessEvents-specific features.

You can debug local projects using their CDD and EAR files, and you can also debug remote engines.



**Declarative Programming and Stepping** Step into, step over, and step return may not behave as you might expect. In a declarative rule language like TIBCO BusinessEvents, unlike with a procedural language, there is no predetermined path through the code. The inference engine logic determines the next action (next rule) based on various conditions and settings, not the rule itself.

### Testing Projects

Running projects to test them works in a way similar to debugging, without the ability to add break points and so on. For testing purposes, you choose options to run the engine, rather than to debug the engine.

## Launch Configurations

Before you debug or test a project, you define settings needed to launch the engine, using a *launch configuration* file.

You can add multiple launch configurations for one project, each configured for a different purpose. For example, you may want to test the effect of different startup arguments. Or if you are using cache OM, you could create one configuration for the processing unit (PU) that runs an inference agent, and one for the PU that runs the cache agent.

## Test Data

When you test or debug an engine, you must assert test data directly to working memory. The test data in the working memory triggers rules in the normal way, so you can observe conflict resolution and run to completion cycles in the engine.

You can send data to a destination as you would in a runtime environment. For example, you can send messages from a JMS or Rendezvous source.

You do not have to have external sources, however. You can also provide test data using the following TIBCO BusinessEvents Studio features:

**Test Data** You can create and save concept and event instance data for later use. You can then assert this data selectively when running or debugging the engine.

**Rule Data** You can provide data as it is expected by rules in the Rule Data tab of the Debugger perspective. See [Working with Rule Data on page 692](#).

See [Creating and Working With Test Data on page 690](#).

## The Rule Agenda and Variable Views

When you run the debugger and reach a breakpoint in a rule or rule function the execution is suspended. You can now examine the state of things:

- The Variable view displays value of all the variables in scope.
- The Rule Agenda view displays rules to be executed from the rule agenda, excluding the current rule (where execution is suspended due to the breakpoint).

For example, Event1 is asserted, which causes Rule1 to execute. Rule1 creates an instance of Concept1. Rule2, Rule3, and Rule4 all have Event1 and Concept1 in scope, and their conditions are satisfied by the instance of Event1 and Concept1. If a breakpoint suspends execution in Rule2, then you see Rule3 and Rule4 in the Rule Agenda view.

For details about how the rule agenda is built, see *Understanding Conflict Resolution and Run to Completion Cycles* in *TIBCO BusinessEvents Architect's Guide*.

## Viewing and Understanding Results

After every run, TIBCO BusinessEvents creates a consolidated results file in XML format with detailed information on the test run. This results XML file is easy to view and compare.

An editor also appears that displays the results in TIBCO BusinessEvents Studio. See [Viewing the Results on page 699](#).



The results are shown for in-memory applications only.

## Debugger Setup

You can create test data and set breakpoints before you run debugger, or using views within the Debug perspective.

Other setup must be done before you begin. See the following sections for details:

- [Preparing to Run \(Test\) or Debug a Project, page 683](#)
- [Adding and Working with Launch \(Debug or Run\) Configurations, page 685](#)
- [Creating and Working With Test Data, page 690](#)
- [Setting Breakpoints in Rules and Rule Functions, page 693](#)

## Preparing to Run (Test) or Debug a Project

---

You can do these tasks in any order. For information on setting up Tester Preferences, see [Tester Preferences on page 735](#).

### Build an EAR File

1. In TIBCO BusinessEvents Studio, select the project in the BusinessEvents Studio Explorer panel, then select **Project > Build Enterprise Archive**.
2. If you want to build an EAR with debug information, check the **Generate Debug Info** checkbox (checked by default).
3. In the File Location field provide the path to the EAR file and specify the EAR file name, for example, `c:\myprojects\myproject.ear`.
4. Generate the EAR file.

### Create Test Data (as Desired)

You can create test data for use across multiple sessions, or you can provide rule input data while the engine is running. See [Creating and Working With Test Data, page 690](#).

### For Remote Debugging Only, Configure Java Debug Interface (JDI)

To configure for remote debugging you configure the `be-engine.tra` file on the remote engine so the engine uses Java Debug Interface (JDI) for remote debugging.

You then configure a debug configuration for remote debugging, as explained in [Adding and Working with Launch \(Debug or Run\) Configurations on page 685](#). You must specify the same JDI port number in the TRA file and in the debug configuration.

#### To Configure Java Debug Interface (JDI)

For each TIBCO BusinessEvents engine you want to enable for remote debugging, do the following.

1. Open the `BE_HOME/bin/be-engine.tra` file for editing.
2. Specify the port on which you want the engine to listen, using the environment variable `tibco.env.JDI_PORT`, for example:

```
tibco.env.JDI_PORT 5192
```

Where 5192 is the default value. If multiple engines run on the same machine, ensure that each has a unique port.

3. Uncomment the following line:

```
-Xrunjdp:transport=dt_socket,address=%JDI_PORT%,suspend=na,server=y
```

4. Start or restart the engine with the `-d` (or `-debug`) command.

## Adding and Working with Launch (Debug or Run) Configurations

Launch configurations are of two types: *run configurations* and *debug configurations*. All launch configurations have the same basic fields, but debug configurations have extra settings for remote debugging.

**Before you begin** Ensure that you know the EAR file location, the CDD file location, and the name of the processing unit (configured in the CDD file) that you want to use. (The processing unit runs as an engine.)

### To Add and Work with Launch Configurations

A reference to the settings is provided in [Launch Configurations Reference on page 688](#).



For remote debugging you first have to set properties in the remote engine TRA file. See [For Remote Debugging Only, Configure Java Debug Interface \(JDI\) on page 683](#).

1. Do one of the following:
  - To create a debug configuration, click the down-arrow to the right of the debugger button (  ) on the toolbar. From the drop-down list, and select **Debug Configurations**. Or, select **Run > Debug Configurations**.
  - To create a run configuration, click the down-arrow to the right of the Run button (  ) on the toolbar. From the drop-down list, select **Run Configurations**. Or, select **Run > Run Configurations**.

You see the Debug Configurations or Run Configurations dialog.
2. Select an option from the list on the left:
  - For testing or local debugging, select **TIBCO BusinessEvents Application**
  - For remote debugging, select **Remote TIBCO BusinessEvents Application**.

3. Do one of the following depending on your needs:
  - To edit a configuration, expand TIBCO BusinessEvents Application or Remote TIBCO BusinessEvents and select an existing debug configuration.
  - To add a new configuration, click the **New Configuration** () button.
  - To duplicate a configuration, select the configuration and then click the **Duplicate** () button. Modify, then save as a new configuration.
  - To delete a configuration, select the configuration and then click the **Delete** () button.

When you add, edit or duplicate a configuration, Configuration fields appear in the right panel.

4. If you selected TIBCO BusinessEvents Application in [step 2](#), do the following:
  - Select the **Main** tab and configure values as explained in [Launch Configurations Reference on page 688](#).
  - Select the **Classpath** tab and configure the classpath for external libraries or custom functions as needed, for example if the project uses Rendezvous or JMS channels. See [Working with External Library and Custom Function Paths on page 8](#) for details.
  - Select the **Environment** tab and configure environment variables as needed, to run or debug the project in TIBCO BusinessEvents Studio. You can add new variables. You can select and then edit existing variables. You can append your edited variable to the existing environment variable, or you can replace the existing environment variable with it. For example if a custom function depends on a native library, you can add the path to that library using the `PATH`, `LD_LIBRARY_PATH`, `SHLIB_PATH`, or `LIBPATH` variable, as appropriate for your operating system.



On Linux platforms when TIBCO BusinessEvents DataGrid is used, you must set `LD_LIBRARY_PATH` to `AS_HOME/lib` in the Environment tab.

- If you need to display characters that are not present in the system's default charset on the console, set the encoding to UTF-8.
- Select the **Common** tab and set the **Encoding** to Other. Then select **UTF-8** from the drop-down menu.

This procedure is documented in Plug-in Development Environment Guide > Reference > Launchers > JUnit Plug-in Test Launcher > Common Tab.

- For information on standard Eclipse features incorporated into this area, see Eclipse help.

5. If you selected Remote TIBCO BusinessEvents Application in [step 2](#), you must connect to the remote engine. Click the **Remote** tab and specify the host name or IP address and port the engine is running on. Use the same JDI port number in the TRA file and in the debug configuration (see [For Remote Debugging Only, Configure Java Debug Interface \(JDI\) on page 683](#)).
6. Click **Apply** to save configuration settings.
7. Do one of the following:
  - Click **Close** and save the configuration you worked on.
  - If you have done the setup as explained in this chapter and are ready to run or debug, click **Debug** to save the configuration and launch the debugger, or click **Run** to save the configuration and run the engine.

You can then assert test data as desired in order to observe the effect of the data on the engine. See [Asserting Rule Input Data on page 697](#).

## Launch Configurations Reference

---

See [Adding and Working with Launch \(Debug or Run\) Configurations on page 685](#) for the related procedure.



The Source and Common tabs are standard Eclipse dialogs. See Eclipse help for details on use of those tabs. If the project uses third party JARs, you must also reference them in the Classpath tab, and update the Environment as needed.

### For Testing and Local Debugging

| Field                | Notes                                                                                                                                                                                                                                                                                               |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                 | A descriptive name. It appears in the drop-down list of configurations.                                                                                                                                                                                                                             |
| <b>Main Tab</b>      |                                                                                                                                                                                                                                                                                                     |
| Project              | Browse to select the name of the TIBCO BusinessEvents project to use for this configuration. You can select from projects in the workspace. The project currently selected in BusinessEvents Studio Explorer appears by default.                                                                    |
| VM Arguments         | Optional. Provide options and parameters using <code>-V</code> , <code>-D</code> , <code>-X</code> and so on.<br>For example to set global variables use: <code>-VVariable=value</code> .                                                                                                           |
| CDD File Location    | Browse to select the CDD file to be used for this launch configuration.                                                                                                                                                                                                                             |
| Processing Unit Name | Select the name of the processing unit (PU) whose values are used for this launch configuration. The drop-down list displays PUs available in the CDD specified in the CDD File Location setting.                                                                                                   |
| Working Directory    | The location of the working directory for the TIBCO BusinessEvents engine. Used to store temporary files and logs. Browse to and select an existing directory.<br><br>Path names that do not start with the root directory are assumed by the operating system to start from the working directory. |
| EAR File             | Browse to select the EAR file to be used for this launch configuration.<br><br>The EAR file must be generated with the Generate Debug Info option checked.                                                                                                                                          |

## For Remote Debugging

You can debug a running TIBCO BusinessEvents engine on the current machine or another machine

| Field             | Notes                                                                                                                                                                                                                                                                                                            |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name              | A descriptive name. It appears in the drop-down list of configurations.                                                                                                                                                                                                                                          |
| <b>Main Tab</b>   |                                                                                                                                                                                                                                                                                                                  |
| Project           | Browse to select the name of the TIBCO BusinessEvents project to use for this configuration. You can select from projects in the workspace. The name of the TIBCO BusinessEvents project in the workspace. It must be the same as the project that is running remotely.                                          |
| <b>Remote Tab</b> |                                                                                                                                                                                                                                                                                                                  |
| Remote Connection | Enter VM arguments for running the remote VM. Alternatively, add these to the <code>java.extended.properties</code> property in the remote application's runtime properties (TRA) file. Default value is:<br><code>-Xdebug</code><br><code>-Xrunjdwp:transport=dt_socket,address=25192,suspend=n,server=y</code> |
| Host              | The host name or IP address of the remote computer where you are running the TIBCO BusinessEvents engine.<br><br>Default is <code>localhost</code> .                                                                                                                                                             |
| Port              | Communication port for the debugger, on the remote machine.<br><br>Default is <code>25192</code>                                                                                                                                                                                                                 |

## Creating and Working With Test Data

---

You can send data to the engine through channels in the normal way. You can also create data within TIBCO BusinessEvents Studio for assertion into the running engine's Rete network during testing and debugging. Doing so means you don't have to have the external resources in place in order to test or debug the runtime.

You can create data while you are working in the debug perspective. You can also create and save test data ahead of time and save it in your project for later use.

The internal data is used at the following *bottom* tabs, within the Rule Input tab:

- **Tester Data** This is concept and event instance data. You create it in the TIBCO BusinessEvents Development perspective, right-click an entity and select Create Test Data. You use it in the Debug perspective: select the Rule Input View > Tester Data tab.
- **Rule Data** This is data directly provided as rule input. It is created and edited in the Rule Input View > Rule Data tab (in the Debug perspective).

How you use test data depends on what aspect of a project you want to test or debug.



If the test data gets asserted in Cache Agent then the studio shows a message "Test data cannot be asserted" along with a "Can't assert Object when activeMode is false" message on console as cache agents are not actively participating in processing inputs. Only Inference Agents processes inputs and the test data should be asserted into Inference Agent.

### Working with Concept and Event Test Data

You can enter data in the Test Data editor for event payloads, and for concept, event and scorecard properties, including properties that are primitive types, array types, contained concepts, and reference concepts. If the concept, event or scorecard properties are associated with a domain model, then the test data gets populated with the values in the domain model. You can use global variables.

#### To Create Concept and Event Instance Test Data

1. In BusinessEvents Studio Explorer (or in the Debug perspective BusinessEvents Studio Explorer view), right-click an event or concept and click **Create Test Data**.

The Test Data editor appears showing the event or concept properties as column headers.

2. In the Test Data editor, click **Add** to add rows for new instances. You can add your own unique **extId** values to the test data input, as needed. You can also use global variables.

You can also remove existing rows by selecting one or more rows and clicking **Remove**.

3. Click **Save**. The entity's test data is saved to an XML file stored within the `TestData` folder in the project root.

The `/TestData` folder is the default location (see [Tester Preferences on page 735](#)).

### To Edit Test Data for Concepts and Events in BusinessEvents Studio Explorer

To add more test data or edit test data you created earlier, do the following.

1. In BusinessEvents Studio Explorer, expand the `TestData` folder in the root of the project.

The `/TestData` folder is the default location (see [Tester Preferences on page 735](#)).

2. Drill down to the test data you want to edit. The folder structure matches the project's event and concept folder structure.
3. Double-click the name of the test data file you want to edit. The test data editor opens. Add, remove, and edit rows of test data as desired.

Test data filenames use the format `eventName.eventtestdata` and `conceptName.concepttestdata`.

### To Edit Concept and Instance Test Data in the Rule Input View

1. Open the Rule Input view, if it is not already shown. To make the Rule Input View visible, do one of the following:
  - Select the Debug perspective as follows: select **Window > Open Perspective**, or click the Open Perspective () button). Then select **Other > Debug**. The views associated with the Debug perspective open.
  - In the **Window** menu, click **Show View > Other**. Expand **TIBCO BusinessEvents**, and select **Rule Input**. Click **OK**.
2. Click the **Rule Input** tab and then select the **Tester Data** bottom tab.

For each concept or event for which you created test data, you see one row showing the project path to that concept or event.

3. Double-click the row for the concept or event whose test data you want to edit.

The Test Data editor appears, showing the rows of test data already created.

4. Edit as desired.

## Working with Rule Data

### To Create and Save Rule Data

1. Open the Rule Input view, if it is not already shown. To make the Rule Input View visible, do one of the following, as needed:
  - Select the Debug perspective: Select **Window > Open Perspective**, or click the Open Perspective () button). Then select **Other > Debug**. The views associated with the Debug perspective open.
  - In the **Window** menu, click **Show View > Other**. Expand **TIBCO BusinessEvents**, and select **Rule Input**. Click **OK**.
2. Provide input from the mapper as explained in [Using the Function Argument Mapper on page 266](#).
3. Specify the Launch Target. This is generally the locally running engine. The other fields become active.
4. Specify an event or concept in the Entity URI field. If you specify an event, then specify the Destination URI. Specify the Rule Session (agent) to use.
5. Do one of the following:
  - Click **Save** to save the input values.  
You can reuse these saved values for repetitive tests.
  - Click **Load** to load the input values from an existing XML file.
  - Click **Assert** to assert the data to a running engine.

## Setting Breakpoints in Rules and Rule Functions

Setting breakpoints is an Eclipse feature. This section provides only basic information. You can also use advanced features such as importing and exporting breakpoints, and using class prepare breakpoints. See Eclipse help for more details on all breakpoint functionality. You can set or change breakpoints during a debug session also.



You cannot step through custom Java code.

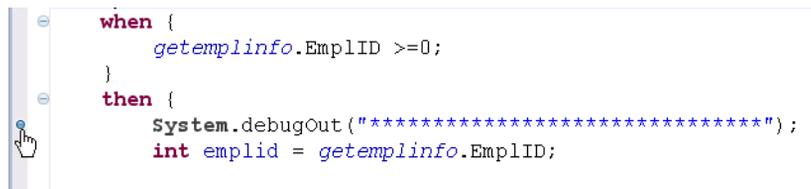
### To Set Breakpoints in Rules and Rule Functions

1. In TIBCO BusinessEvents Studio, open the source editor for a rule or rule function.

You can work with breakpoints in the debugger perspective as well as in the TIBCO BusinessEvents Studio development perspective.

2. To add a breakpoint put your cursor in the left margin (gray area) next to a row where you want to add a breakpoint. Do one of the following:
  - Right click and select Toggle Breakpoint.
  - Double-click in the left margin.

A breakpoint icon appears in the left margin:



3. Set and adjust breakpoints as needed. Select a break point, right click, and do any of the following:
  - To disable a breakpoint, select Disable.
  - To remove a breakpoint, select Toggle Breakpoint. (Or just double click the breakpoint.)
  - To edit a breakpoint's properties, select Breakpoint properties. A dialog displays (with mostly runtime options). For example, you can use a class prepare breakpoint (so the running program is suspended when the specified class or interface is first loaded by the Java VM).

A breakpoint may not be set exactly where you place it. This is because TIBCO BusinessEvents ensures that breakpoints fall on an executable statement, and moves any that do not to the nearest executable statement.



**Executing step code line does not match rule editor code line** This situation happens when there is a mismatch between the debug line information stored in the EAR, and the information in the open rule editor. To resolve the problem, recompile the EAR file.

## Running Debugger

Make sure you have completed setup as needed. Then run debugger as explained below.



For projects containing an RV channel, add the path to the location of the RV libraries by adding `RV_HOME/lib` to the existing path.

### To Run Debugger

1. As needed, switch to Debug perspective. Select **Window > Open Perspective**, or click the Open Perspective () button). Then select **Other > Debug**.

Alternatively, wait till TIBCO BusinessEvents prompts you to change to debug perspective. This happens when the debugger reaches the first breakpoint.

2. Click the down-arrow to the right of the debugger ( ▾) button. You see a drop-down list. Do one of the following:
  - Select a debug configuration from the list.



To add configurations to the drop-down list, select Organize Favorites from the list.

- Select Debug Configurations. At the Debug Configurations dialog select a debug configuration and click **Debug**.
- Click the debugger () button or Select **Run > Debug**. (Only if you have already launched debugger with a configuration.)

Debugger starts a TIBCO BusinessEvents engine, using parameters provided in the launch configuration, if any were provided.

3. Assert data as explained in [Asserting Rule Input Data](#).
4. View the results as explained in [Viewing the Results on page 699](#).
5. Use the standard Eclipse commands such as step into (F5), step over (F6), step return, step return, and so on, depending on what you want to examine. See the options in the Run menu and in the Breakpoints tab for more options, and use Eclipse help for details with these Eclipse features.
6. When you are finished, click the Terminate () button to stop the engine.

## Running Tester

---

Make sure you have completed setup as needed. You can create test data before you run tester, or using views within the Debug perspective. Other setup must be done before you begin.



For projects containing an RV channel, add the path to the location of RV libs:

See the following sections for details:

First do the setup needed for your testing. See the following sections:

- [Preparing to Run \(Test\) or Debug a Project, page 683](#)
- [Adding and Working with Launch \(Debug or Run\) Configurations, page 685](#)
- [Creating and Working With Test Data, page 690](#)

Then run tester as explained below.

### To Run Tester

1. As needed, switch to Debug perspective. Select **Window > Open Perspective**, or click the Open Perspective () button). Then select **Other > Debug**.
2. Select **Run > Run Configurations** or click the down-arrow to the right of the Run () button and choose **Run Configurations**.

If you have configured favorites, you can click the down-arrow to the right of the Run () button and choose a favorite.

3. At the Run Configurations dialog select a run configuration and click **Run**.



If you have already started an engine using a run configuration and want to start it again, click the Run () button or Select **Run > Run**.

A TIBCO BusinessEvents engine starts, using parameters provided in the run configuration, if any were provided.

4. Assert data as explained in [Asserting Rule Input Data](#).
5. View the results as explained in [Viewing the Results on page 699](#).
6. When you are finished, click the Terminate () button to stop the engine.

## Asserting Rule Input Data

---

You can assert either tester data, or rule data for use in Tester or Debugger.



To perform any task in the **Rule Input** tab, you must keep the engine running. You can also send messages to destinations, as you would at runtime.

See also [Creating and Working With Test Data on page 690](#) for details about creating and saving data.



Test data you created earlier for concepts, events, or scorecards appears in the Debugger perspective. By default the view appears in the middle on the left.



For cache application usign preprocessors to load a concept from cache, if an event does not have the default destination then that destination has to be assigned to event before starting test in Rule Input View

### To Assert Tester Data

1. Select the Rule Input tab and then the Tester Data bottom tab.
2. One entity can have multiple rows of test data. To select which row or rows of test data to assert, double click the entity's URI in the Select Test Data view. It appears in the Test Data editor. Check one or more checkboxes in the Use column as desired. You can repeat this step for all the entities shown in the Select Test Data view, as needed.

You can edit the test data at this time too and you can add test data for more entities. See [Working with Concept and Event Test Data on page 690](#) for details.

3. In the Input panel, Launch Target field, specify which engine to use.
 

Multiple engines can run in tester at the same time, for example, a cache agent engine, and an engine running inference agents.
4. In the Input panel, Rule Session field, specify which agent to assert the data to.
 

One engine (processing unit) can have multiple inference agents.
5. Click **Start Test**.

You see console messages and results. See [Viewing the Results on page 699](#) for details about understanding the results of a test or debugger run.

6. As appropriate, select more test data to assert (as in [step 2](#)), and again click **Start Test**. Once again, analyze the results of asserting that data using the other views.

### To Run Tester or Debugger with Rule Data

1. Open the Rule Input view, if it is not already shown. In the **Window** menu, click **Show View > Other**. Expand **TIBCO BusinessEvents**, and select **Rule Input**. Click **OK**.
2. Do one of the following:
  - Provide input from the mapper as explained in [Using the Function Argument Mapper on page 266](#). Optionally you can click **Save** and save the values to an XML file.
  - Click **Load** to load the input values from an existing XML file.
3. Specify the **Launch Target**, **Entity URI**, **Destination URI** and **Rule Session**.
4. Click **Assert**.

See [Viewing the Results on page 699](#) for details about understanding the results of a test or debugger run.

## Viewing the Results

After data is asserted to the working memory (see [Running Tester on page 696](#)), a run to completion (RTC) cycle occurs and you see the following:

- The Console tab displays engine console messages.
- The result data editor appears showing the results of the run. The editor title displays the result data filename with the format `Run-n.resultdata`.
- The results of this test are stored in an XML format in a `.resultdata` file

By default the results are stored in the `/TestData/Project Name/Processing Unit Name` folder. See [Tester Preferences on page 735](#) for information on changing the location where test data is stored.

You can also open the result data in its editor by double-clicking the `.resultdata` file in BusinessEvents Studio Explorer.

The screenshot displays the BusinessEvents Studio Explorer interface for a test run titled "Test Result: FraudDetection/default/Run-1".

**Created:**

- /Concepts/Account{id = 5}
- /Rules/ProcessDebits/CreateAccount [Invoked]
  - /Concepts/Account{id = 4} [Causal]

**Modified:**

- /Concepts/Account{id = 8}
- /Rules/ProcessDebits/ApplyDebit [Invoked]
  - /Concepts/Account{id = 8} [Causal]
  - /Events/Debit{id = 14} [Causal]

**Deleted:**

- /Events/Debit{id = 14}

**Result Test Data:**

**After**

| Account @id | Account @ExtId | Balance | Debits | Status | AvgMonthlyBal... |
|-------------|----------------|---------|--------|--------|------------------|
| 8           | ActB           | 27000.0 | 3000.0 | Normal | 15000.0          |

**Before**

| Account @id | Account @ExtId | Balance | Debits | Status | AvgMonthlyBal... |
|-------------|----------------|---------|--------|--------|------------------|
| 8           | ActB           | 30000.0 | 0.0    | Normal | 15000.0          |

## Understanding Result Data

The Result Data view shows the results of a single RTC cycle. On the left of the view you see the **Created**, **Modified**, and **Deleted** entities (if any).

Expanding each entity shows the rule that affected the entity, and also the causal object that triggered the rule.

Click any entity to display its values in the **Result** panel.



Test results will not be shown for cache applications; they will be shown only for in memory applications.

## Viewing and Understanding Working Memory Contents

Option to modify Working Memory Contents is available only in the Working Memory editor.

The Working Memory concepts or scorecards can be manipulated in Working Memory View. Change the property values and click modify.

The screenshot displays the Studio Explorer on the left, showing a project structure for 'CreditCardApplication' with folders like Channels, Concepts, Domains, Events, Rule\_Functions, Rule\_Templates, Rules, TestData, Virtual\_RF, and CC.cdd. The main window shows the 'Working Memory Contents:inference-class' view, which lists several objects with their IDs. The 'Result Test Data' panel on the right shows a table with the following data:

| CreditCardApplication @ExtId | CreditCardApplication @ExtId | Name        | CreditType | CardIssued                           |
|------------------------------|------------------------------|-------------|------------|--------------------------------------|
| 5                            | application2                 | Sandra Weir | Standard   | be.gen.Concepts.CreditCard@id=14&v=0 |

### To see the Working Memory Contents

1. Click the Working Memory Contents icon on the tool bar, or from the **Project** menu.
2. Select the appropriate rule session (agent).

The number of objects in the working memory contents is specified in the Preferences. For more details, see [Tester Preferences on page 735](#).

## Chapter 41 **Diagrams**

This chapter covers TIBCO BusinessEvents diagrams, including a summary of the diagrams available, how to access and work them, and what options and preferences are available for working with them.

### Topics

---

- [Overview of Diagrams, page 702](#)
- [Working with Diagrams, page 704](#)
- [Project Analyzer and Selected Entity Project Diagrams, page 712](#)
- [Dependency Diagrams, page 715](#)
- [Sequence Diagrams, page 717](#)
- [Concept Model Diagrams, page 718](#)
- [Event Model Diagrams, page 719](#)
- [Diagram Options and Tools Reference, page 720](#)

## Overview of Diagrams

---

TIBCO BusinessEvents provides several kinds of diagrams, which are visualization tools that help you to understand and analyze even very large and complex projects.



The diagrams in TIBCO BusinessEvents are based on Unified Modeling Language (UML), but are NOT completely compliant to UML.

The diagrams allow you to show or hide details. You can open editors for any project component in a diagram. You can also create snapshot images of diagrams to share project information with other personnel.

The main types of diagrams, and the project elements that use them, are as follows:

**Selected Entity Project diagrams** These diagrams display all the selected project resources, and how they interact with one another. External resources such as XSD and WSDL files are not shown. You can either create a project diagram for an entire project, or only for the selected resources of a project. You can also choose to run the project analyzer while creating a project diagram, which analyzes the project resources and gives a report in the Problems view.

**Dependency diagrams** These diagrams show the relationships between the selected project resource and its dependent resources. They are available for channels, concepts, state models, scorecards, events (all types except advisory events), rule functions, and rules.

**Sequence diagrams** These diagrams show how project resources are called into use at runtime. They are available for events (all types except advisory events), rule functions, and rules.

**Concept model diagrams** These diagrams show the concept model for a project. You can view a model diagram for all the concepts in a project.

**Event model diagrams** These diagrams show the event model for a project. You can view a model diagram for all the events in a project, except for advisory events.



**Add-on Products** State model diagrams are available in the TIBCO BusinessEvents Data Modeling add-on. Metrics are available in the TIBCO BusinessEvents Views add-on. For metrics, you can view a dependency diagram only by right-clicking in Project View. See the add-on product documentation for more details

Table 68 Types of Diagrams Available for Different Resource Types

| Project Resource | Project | Dependency | Sequence | Concept Model | Event Model |
|------------------|---------|------------|----------|---------------|-------------|
| Channel          | Yes     | Yes        |          |               |             |
| Concept          | Yes     | Yes        |          | Yes           |             |
| Domain           | Yes     | Yes        |          |               |             |
| State Model      | Yes     | Yes        |          |               |             |
| Scorecard        | Yes     | Yes        |          |               |             |
| Event            | Yes     | Yes        | Yes      |               | Yes         |
| Rule Function    | Yes     | Yes        | Yes      |               |             |
| Rule             | Yes     | Yes        | Yes      |               |             |
| Shared Resource  | Yes     | Yes        |          |               |             |
| Metric           | Yes     | Yes        |          |               |             |

## Working with Diagrams

---

You can work with diagrams in a variety of ways for your own information, and you can use them to share aspects of your project with others in a visual way that is easy to understand. This section presents some common tasks and lists the various tools you can use to work with all types of diagrams.

### Configuring Diagram Preferences and Properties

#### Diagram Preferences

See [Diagram Preferences on page 730](#) to understand and set preferences so that diagrams display as you desire.

#### Diagram Properties

Use the following properties to handle diagrams that are too complex to display in a timely manner. You can adjust the values in the `BE_HOME/studio/eclipse/configuration/studio.tra` file to define upper limits:

---

```
#Diagram properties
be.studio.project.diagram.nrEdges=3000
be.studio.project.diagram.nrMaxEdges=1500
```

---

If the number of edges in the diagram exceeds the limit defined by `be.studio.project.diagram.nrEdges`, then the usage links in the diagram are hidden. Usage links are links from a rule or rule function to other rules or rule functions it uses.

If the number of edges in the diagram exceeds the limit defined by `be.studio.project.diagram.nrMaxEdges`, then symmetric layout is run as default rather than hierarchical because symmetric layout is faster and less expensive.

## Different Ways to Create Diagrams

You can create diagrams in various ways as shown in the table below.

| Method of creating diagrams                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Diagrams you can create                                                                         |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| In BusinessEvents Studio Explorer, right-click a resource or a project, and select the appropriate menu option                                                                                                                                                                                                                                                                                                                                                                                                       | Selected Entity Project Diagram, Concept Model Diagram, Event Model Diagram                     |
| In Studio Explorer, select a resource and right-click                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Selected Entity Project Diagram, Concept Model Diagram, Event Model Diagram, Dependency diagram |
| In a resource editor, click the appropriate diagram button in the top right corner<br>Concept Model Diagram <br>Event Model Diagram <br>Dependency Diagram <br>Sequence diagram  | Concept Model Diagram, Event Model Diagram, Dependency Diagram, Sequence Diagram                |
| In a Selected Entity Project diagram, right-click a resource                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Dependency Diagram                                                                              |
| In a Selected Entity Project diagram, right-click anywhere on the canvas                                                                                                                                                                                                                                                                                                                                                                                                                                             | Selected Entity Project diagram                                                                 |

## Performing Common Tasks

You can perform the following common tasks on all types of diagrams:

- Hover the mouse over a resource to view a tooltip with details
- Double-click any resource represented in the diagram to open its editor
- Drag resources to any location in order to change the layout as desired
- Drag any link to change the layout as desired
- Click any arrowhead to move the arrow to any other element in the project. When you hover the mouse over an arrowhead, you see a small graphic on it.



- Double-click the black and white plus signs in a dependency or sequence diagram to unfold the next level of dependency

## Using Diagram Tools

You can work with diagrams with the help of the following tools available from the Diagram menu:

- [Interactive Tools](#), such as Select, Pan, Magnify, Marquee Zoom, Interactive Zoom, Link Navigator
- [Layout](#), such as Default Layout, Circular Layout, Orthogonal Layout, Symmetric Layout, Rectilinear Hierarchical Layout, and Oblique Hierarchical Layout
- Print tools: Print Setup, Print Preview, Print
- Export to Image
- Labeling
- Link Routing
- Incremental Layout
- Fit In Window

The above tools are also available on the toolbar when you create or display a diagram. The toolbar also has three additional tools listed below:

- Refresh Diagram
- Zoom Percentage
- Search Diagram Entities

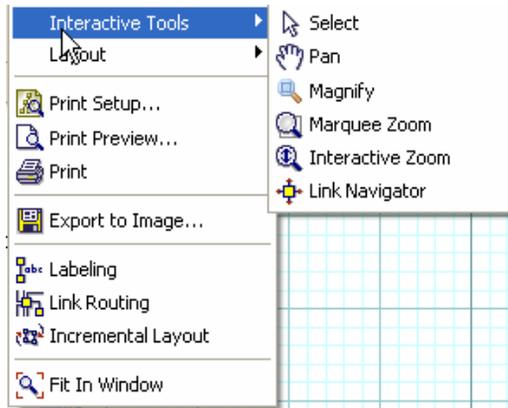


Diagram menu and the toolbar described in this section are available only when you create or open a diagram.

For more details about the diagram tools, see [Diagram Options and Tools Reference on page 720](#).

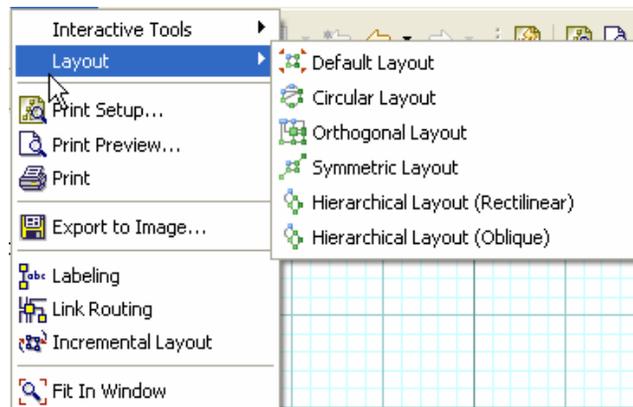
## Interactive Tools

To access the Interactive Tools, select Interactive Tools from the Diagram menu. They are also available on the toolbar when you create or display a diagram. Interactive Tools let you select, pan, and magnify diagrams. They also include tools such as marquee zoom, interactive zoom, and link navigator. For more details about these tools, see [Diagram Options and Tools Reference on page 720](#).



## Layout

You can access the Layout menu from the Diagram menu. They are also available on the toolbar when you create or display a diagram. Layout options let you change the layout of a diagram to Circular, Orthogonal, Symmetric, or Hierarchical. For more details about these tools, see [Diagram Options and Tools Reference on page 720](#).



## Context Menu Diagram Tools

Diagram tools are available as context menus for diagram canvas and objects. [Table 69, Context Menu Options for Canvas and Objects](#) lists these context menu options.

*Table 69 Context Menu Options for Canvas and Objects*

| Context Menu Option             | Available for                            |
|---------------------------------|------------------------------------------|
| Selected Entity Project Diagram | Canvas – Selected Entity Project Diagram |
| Export to Image...              | Canvas – All diagrams                    |
| Print Setup...                  | Canvas – All diagrams                    |
| Print Preview...                | Canvas – All diagrams                    |
| Fit In Window                   | Canvas – All diagrams                    |
| Create Dependency Diagram       | Object – Selected Entity Project Diagram |
| Fold One Level                  | Object – Selected Entity Project Diagram |
| Fold N Levels...                | Object – Selected Entity Project Diagram |
| Fold All Levels                 | Object – Selected Entity Project Diagram |
| Edit                            | Object – All diagrams                    |

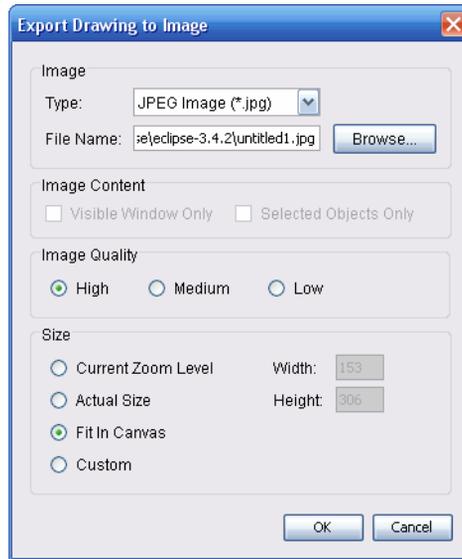
For more details about the diagram tools, see [Diagram Options and Tools Reference on page 720](#).

## Exporting a Diagram to an Image

The Export to Image option enables you to save the diagrams locally for later use.

## To Export a Drawing (Diagram) to Image

1. From the top menu, select **Diagram > Export to Image**. You see the Export Drawing to Image dialog.



2. Select the type of the image from the Type list.
3. In the File Name field, specify the file name along with the location.
4. Check the check boxes **Visible Window Only** and **Selected Objects Only** as appropriate to specify the content of the image.



Selected Objects Only gets enabled only if you select an object in the diagram.

5. Specify the Image Quality as appropriate.
6. Specify the size, and click OK.

The diagram is exported to an image, and saved on the system.

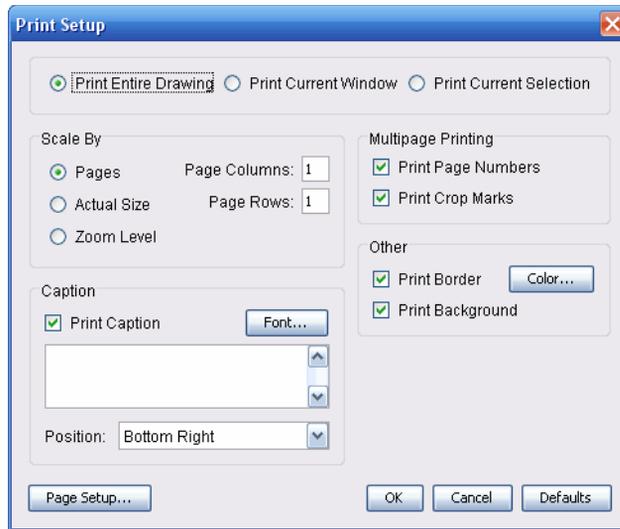
## Printing a Diagram

You can set up preferences, preview the print job, and then print the diagram.

### To Set up Printing Preferences

The Print Setup dialog enables you to specify the print setup for a diagram.

1. Display the diagram, and from the top menu select **Diagram > Print Setup**.



2. Select an appropriate option to print the entire diagram, or to print only the current window, or to print the selection.
3. In the **Scale By** section, select an option to scale the diagram by pages, by actual size or by zoom level.
4. If you want multiple prints, you can specify if you want the page numbers and crop marks to be printed, by clicking to select the respective checkboxes.
5. If you want a caption for your diagram, you can type a caption in the text box that appears on enabling the **Print Caption** checkbox. You can also select a Font for this caption, and its position on the page.
6. If you want to have borders for your diagrams, click to select the **Print Border** checkbox. Click **Color...** to specify the color of the border.
7. If you want to print the background of the diagram, click to select the **Print Background** checkbox.
8. Use **Page Setup** to specify the size, orientation, and other page properties.
9. If you want to reset all the options to the defaults, click **Defaults**.
10. Click **OK**.

### To Preview and Print a Diagram

1. From the top menu, select **Diagram > Print Preview**. An image of the diagram to be printed displays.
2. From the top menu select **Diagram > Print**.

## Setting Diagram Preferences

You can set the preferences for all types of diagrams.

### To Set Diagram Preferences

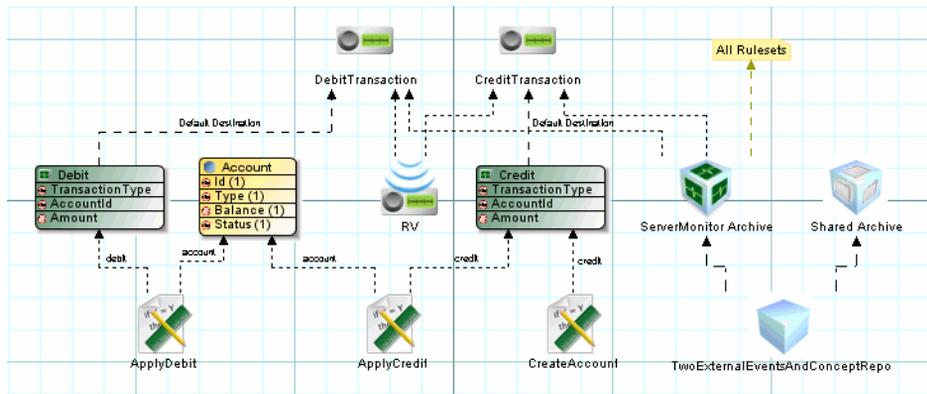
For more information on Preferences, see [Chapter 42, TIBCO BusinessEvents Studio Preferences, on page 725](#).

1. From the top menu, select **Window > Preferences**.
2. In the Preferences dialog, expand **TIBCO BusinessEvents**, and then expand **Diagram**.
3. Select any type of diagram.  
On the right side, you see preferences for the selected diagram type.
4. Edit the preferences, and click **Apply**.

OR

To set all preferences to their default values, click **Restore Defaults**.

## Project Analyzer and Selected Entity Project Diagrams



Project Analyzer and the Selected Entity Project diagram work together to provide insights into your project. Project Analyzer is a document generated in the Problems view area. Selected Entity Project diagram provides a visualization of the whole project or of selected project elements, including dependencies. Selected Entity Project diagrams can display all the resources in a project, how they interact with one another, and how they use or update other resources at run time. You can see concepts, events, rules, rule functions, channels, destinations, score cards, shared resources, and all other resources of a particular project in a Selected Entity Project diagram. It also shows all dependencies in a project.

You can set preferences to determine whether project analyzer runs when you create a Selected Entity Project diagram or not. You can use a menu option to create the project analyzer report without creating the Selected Entity Project diagram.

### Types of Links

A Selected Entity Project diagram has three types of links:

- Links, represented by continuous lines, indicate inheritance or containment
- Scope Links, represented by dashed lines, indicate resources in the scope
- Usage Links, represented by dashed lines with dots, indicate resources that can be used

## Advantages of Project Analyzer and Selected Entity Project Diagrams

Using Project Analyzer and Selected Entity Project Diagrams together helps you understand your project in these ways:

- Helps you understand how all the selected project resources are connected and how they interact
- Analyzes rules and how they modify ontology
- Lets you search entities in a diagram
- Lets you choose the entities to be displayed in a diagram by using Project Filters
- Displays statistics of project resources (number of concepts, events, state machines, and so on) in the properties view
- Enables you to share project design, logic, deployments with others in a format they can read without having to have TIBCO BusinessEvents Studio, as well as print the entire project according to page setup

The Project Analyzer report helps you analyze the project by performing the following tasks:

- Finds rules that can never be fired
- Finds rule functions that are never used
- Finds events that are never used/may never be fired
- Finds destinations with no default events
- Finds domain models that are never used, or that are not associated with entity properties
- Finds state models that are orphaned, or that are not associated with any concepts

## Working with Project Analyzer and Selected Entity Project Diagrams

You can show the diagram for an entire project or for selected elements only. You can also filter the diagram to show only certain types of project element.



You can choose whether to run Project Analyzer whenever you create a Selected Entity Project diagram. To do this, from the **Window** menu, select **Preferences > TIBCO BusinessEvents > Diagrams > Project**. Check the **Run Analysis When Creating View** checkbox, on the right.

You can also run Project Analyzer separately in either case. See [Project Analyzer and Selected Entity Project Diagrams, page 712](#).

### To Run Project Analyzer Without Selected Entity Project Diagram

Right-click the top level project folder, and click **Analyze**. Project Analyzer displays the report in the Problems view.

### To Create a Selected Entity Project Diagram

1. Do one of the following to create a Selected Entity Project diagram for the entire project:
    - In BusinessEvents Studio Explorer, right-click a project resource and select **Create Selected Entity Project Diagram**.
    - Double-click the \*.projectview file in the Studio Explorer.
- OR, to create a Selected Entity Project diagram for selected resources
- In Studio Explorer, select resources of interest, right-click, and select **Create Selected Entity Project Diagram**.



When you select a folder all elements in that folder are selected.

The Selected Entity Project diagram appears in the editor area. The diagram tab label reflects the project name and the type of diagram: *projectName.projectview*, for example, *ExternalEventRepo.projectview*.

By default, the Selected Entity Project diagram displays the selected objects and their immediate dependencies. You can change the depth of dependencies shown using preferences.

### To Filter a Selected Entity Project Diagram

To filter what you see, check and uncheck the options in the Project Filter list within the Palette view, and click **Apply**.

Options show or hide elements of the specified type, or group them:

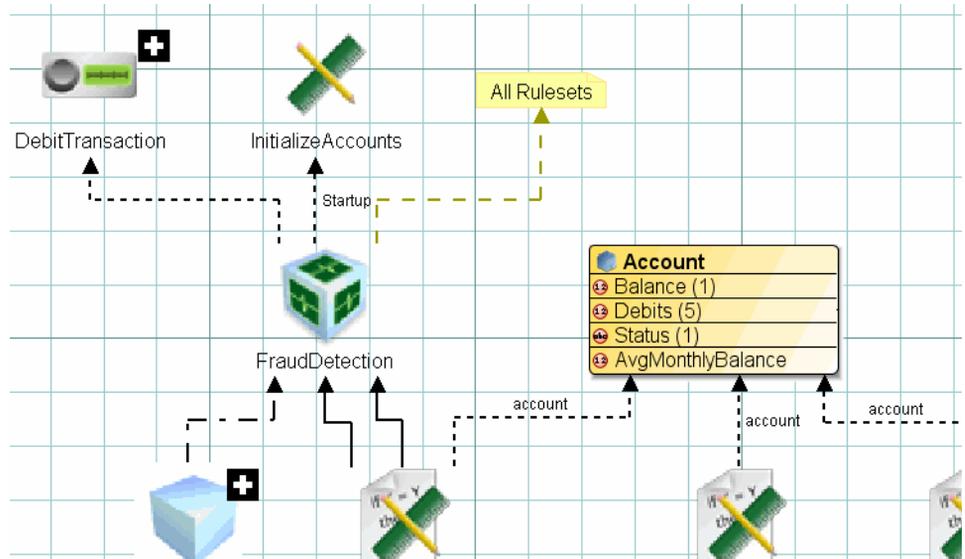
**Show Options** Show Concepts, Show Events, Show Decision Tables, Show Domain Model, Show State Models, Show Archives, Show Rules, Show Rule Functions, Show Scorecard, Show Channels, Show Scope Links, Show Usage Links, Show Archived Destinations, Show Archived Rules, Show Archived Rules (All), Show Rules in Folders, Show Tooltips. If add-on products are used additional options may appear.

**Group Options** Group Concepts, Group Events, Group Rules, Group Rule Functions. The filter options in the Palette view are loaded from diagram preferences. See [Diagram Preferences on page 730](#) for more details.

## Dependency Diagrams

A dependency diagram shows the dependencies for a selected project resource or for the entire project (Selected Entity Project Diagram).

When two resources are dependent and one changes, it can cause changes to the other resource as well. You can control how many levels of dependencies to view in a dependency diagram - one, two, or all levels.



Rule actions are not represented in rule dependency diagrams. Use Debugger to analyze rule actions. See [Chapter 40, Testing and Debugging Projects](#), on page 679.

### To Create a Dependency Diagram

1. Do one of the following:
  - In BusinessEvents Studio Explorer, right-click a project resource and select **Create Dependency Diagram**.
  - Open the project element for editing and click the Dependency Diagram (  ) button in the top right of the editor.
  - In a Selected Entity Project diagram, right-click a resource and select **Create Dependency Diagram**.

The dependency diagram appears in the editor area. By default, one level of dependency is shown (showing only direct dependents). White crosses

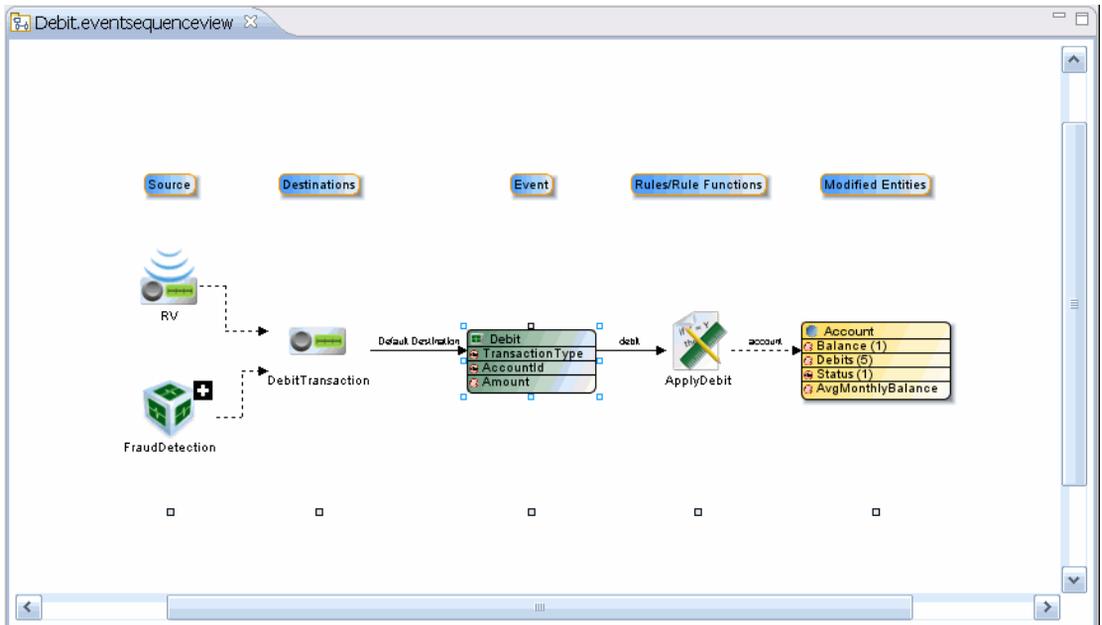
indicate hidden dependencies. The default appearance is also driven by the diagram preferences. See [Diagram Preferences on page 730](#) for more details.

The diagram tab label reflects the item being graphed and the type of diagram: *itemName.item-typedependencyview*, for example, `Transaction.eventdependencyview`.

2. As desired, select a dependency level from the palette:
  - One: Shows the direct dependencies of the selected item
  - Two: Shows the direct dependencies of the selected item as well as the direct dependencies of its direct dependencies
  - All: Shows all dependency relationships that involve the selected item

## Sequence Diagrams

Sequence diagrams capture the behavior of objects and the messages that are passed between them. You can view sequence diagrams for rule functions, rules, and events.



### To Create a Sequence Diagram

Open an event, rule function, or rule for editing and click the Sequence diagram (  ) button in the top right of the editor.

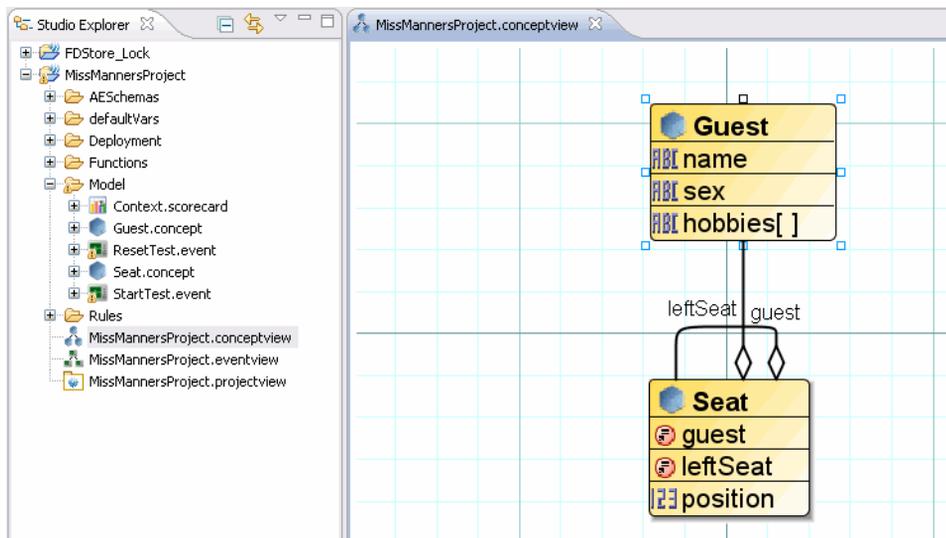
The sequence diagram for the item appears in the editor area.

The diagram tab label reflects the item being graphed and the type of diagram: *itemName.item-typesequenceview*. For example, *Transaction.eventsequenceview*.

## Concept Model Diagrams

The concept model diagrams show inheritance, containment, and reference relationships between all concepts in a project.

Concept model diagrams are created fresh each time you view them (they do not persist on disk).



### To Create a Concept Model Diagram

Do one of the following:

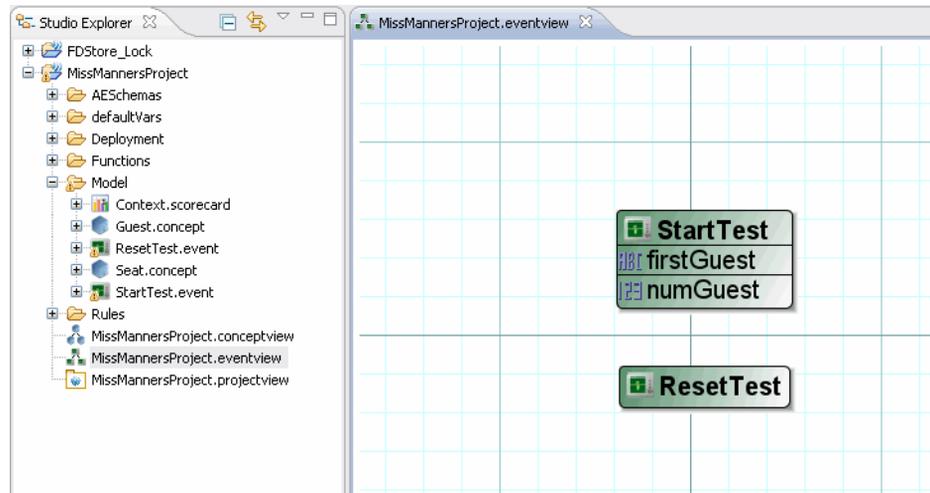
- Right-click a resource or a project in BusinessEvents Studio Explorer and select **Create Concept Model Diagram**
- Open a concept editor and click the Concept Model Diagram (  ) button in the top right of the editor
- Double-click the \* .conceptview file in the Studio Explorer.

The concept model diagram for that project appears, showing only its direct dependents (Dependency level One).

## Event Model Diagrams

The event model diagrams show inheritance relationships between all types of events, except for advisory events.

Event model diagrams are created fresh each time you view them (they do not persist on disk).



### To Create an Event Model Diagram

Do one of the following:

- Right-click a resource or a project in BusinessEvents Studio Explorer and select **Create Event Model Diagram**.
- Open an event editor and click the Event Model Diagram (  ) button in the top right of the editor. The event model diagram for that project appears, showing only its direct dependents (Dependency level One).
- Double-click the \*.eventview file in the Studio Explorer.

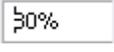
## Diagram Options and Tools Reference

Most diagram options and tools are self-explanatory, and are just shown in the images. This section provides details for a few that require some explanation.

Table 70 TIBCO BusinessEvents Diagram Tools Reference

| Button                                                                              | Description                                                                                                                                                                       |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interactive Tools Menu Options and Toolbar Buttons</b>                           |                                                                                                                                                                                   |
|    | <b>Select</b> (Pointer) tool.                                                                                                                                                     |
|    | <b>Pan</b> tool.                                                                                                                                                                  |
|    | <b>Magnify</b> Select Magnify and hover over the diagram to display a magnified section of the diagram.                                                                           |
|    | <b>Zoom</b> Drag the cursor diagonally to define the area to which you want to zoom.                                                                                              |
|    | <b>Interactive Zoom</b> Drag the cursor up to zoom out. Drag the cursor down to zoom in.                                                                                          |
|    | <b>Link Navigator</b> Shows the path of a link using a simple animation.                                                                                                          |
| <b>Layout Menu Options and Toolbar Buttons</b>                                      |                                                                                                                                                                                   |
|  | <b>Default Layout</b> Returns the diagram to the default layout. To display the layout list, click the arrow next to the button and select a new layout from the list of options. |
|  | <b>Circular Layout</b> It is available in the layout list. See <a href="#">Circular Layout on page 722</a> .                                                                      |
|  | <b>Orthogonal Layout</b> It is available in the layout list. See <a href="#">Orthogonal Layout on page 722</a>                                                                    |
|  | <b>Symmetric Layout</b> It is available in the layout list. See <a href="#">Symmetric Layout on page 722</a> .                                                                    |
|  | <b>Hierarchical Layout (Rectilinear)</b> It is available in the layout list. See <a href="#">Hierarchical Layout on page 723</a> .                                                |

Table 70 TIBCO BusinessEvents Diagram Tools Reference

| Button                                                                              | Description                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | <b>Hierarchical Layout (Oblique)</b> It is available in the layout list. See <a href="#">Hierarchical Layout on page 723</a> .                                                                                                                                                                                                     |
| <b>Other Diagram Menu Options and Toolbar Buttons</b>                               |                                                                                                                                                                                                                                                                                                                                    |
|    | <b>Print Setup.</b> Sets options for printing the diagram or elements of the diagram. See <a href="#">Printing a Diagram on page 709</a> .                                                                                                                                                                                         |
|    | <b>Print Preview.</b> Lets you see how the printed diagram would look.                                                                                                                                                                                                                                                             |
|    | <b>Print.</b> Prints the diagram.                                                                                                                                                                                                                                                                                                  |
|    | <b>Export to Image.</b> Exports the diagram to one of the five available image file formats, and lets you save the image file locally for future reference. See <a href="#">Exporting a Diagram to an Image on page 708</a> .                                                                                                      |
|    | <b>Labeling</b> Prevents labels from overlapping with other elements and labels in the diagram.                                                                                                                                                                                                                                    |
|    | <b>Link Routing</b> Redraws the diagram, attempting to redraw only the links, leaving resource nodes in the same size and position where possible.<br><br>The behavior of link routing is affected by the preference options Fix Node Size and Fix Node Position. See <a href="#">Diagram Preferences on page 730</a> for details. |
|  | <b>Incremental Layout</b> See <a href="#">Incremental Layout on page 723</a> .                                                                                                                                                                                                                                                     |
|  | <b>Fit In Window</b> Zooms the diagram to fit the size of the current diagram editor view.                                                                                                                                                                                                                                         |
| <b>Additional Toolbar Buttons</b>                                                   |                                                                                                                                                                                                                                                                                                                                    |
|  | <b>Refresh Diagram.</b> If you make changes in the project, clicking Refresh Diagram updates the diagram with the changes.                                                                                                                                                                                                         |
|  | <b>Zoom Percentage</b> Zooms the diagram to the specified percentage.                                                                                                                                                                                                                                                              |
|  | <b>Search Entities.</b> Allows you to search for TIBCO BusinessEvents entities in the diagram.                                                                                                                                                                                                                                     |

## Layout Options

Layout options refer to the style of the diagram that is rendered in the diagram editor, Rule Debugger, and Dependency panel.



### Additional Layout Options

Various options affect the appearance of a diagram, no matter what layout is chosen. (The hierarchical layout offers additional options.) See [Diagram Preferences on page 730](#) for more details.

The main layout options are:

- [Circular Layout, page 722](#)
- [Orthogonal Layout, page 722](#)
- [Symmetric Layout, page 722](#)
- [Hierarchical Layout, page 723](#) (default)
- [Incremental Layout, page 723](#)

### Circular Layout

This layout is useful for ontologies where nodes tend to have a clustered (ring or star) structure (where each main node has a starburst of related nodes).

### Orthogonal Layout

The orthogonal layout style uses only horizontal and vertical links. Diagrams are drawn quickly.

The algorithm places highly connected nodes closer together, resulting in a more compact diagram. Even when lines overlap, the algorithm ensures that they are still easy to follow. Because this style has no hierarchical or other visual constraints, the resulting diagrams are often very clear.

### Symmetric Layout

The symmetric layout style looks for and emphasizes the symmetries in a project topology. It can produce a pleasing visual result, if there are no reasons to arrange the nodes. For example, there is no hierarchy or ring-clustering inherent in the structure of the nodes.

See [Configuring Diagram Preferences and Properties on page 704](#) for a property that affects whether symmetric or hierarchical layout is used as the default, depending on number of edges.

## Hierarchical Layout

The hierarchical layout style indicates dependencies by positioning the nodes at different levels. The hierarchical layout style is useful when you need to show precedence relationships in the ontology.

Hierarchical layout is the default. However, see [Configuring Diagram Preferences and Properties on page 704](#) for a property that affects whether symmetric or hierarchical layout is used as the default, depending on number of edges. You can use Layout Preferences to determine in the direction of the hierarchy.

Preferences and options specific to the hierarchical layout are:

- Orientation Options (Edit > Preferences)—Left to Right, Top to Bottom, Right to Left, Bottom to Top
- Routing Options (View > Layout)—Orthogonal or normal (polyline) routing
- Routing Options (Edit > Preferences)—Orthogonal or polyline routing.

See [Diagram Preferences on page 730](#) for more details.

## Incremental Layout

Choose Incremental Layout if you want the main arrangement of the diagram to remain stable when you make changes to your project and only re-route the entities and links that have changed since the last rendering.

Keeping most things in the same place makes it easier to see how changes you have made in the project affect the diagram.

If you want to refresh the entire diagram, click the desired layout option. The program has greater freedom to optimize the layout. The result is likely to be a more pleasing arrangement.

To perform an incremental layout update, do one of the following:

- Select View > Layout > Incremental Layout
- Click the Incremental Layout button  on the toolbar



## Chapter 42 **TIBCO BusinessEvents Studio Preferences**

This chapter shows the preferences available in TIBCO BusinessEvents Studio.

### Topics

---

- [Setting Preferences, page 726](#)
- [Decision Table Related Preferences, page 728](#)
- [Diagram Preferences, page 730](#)
- [Tester Preferences, page 735](#)

## Setting Preferences

---

To access preferences select **Window > Preferences** and expand **TIBCO BusinessEvents**. Also see [Setting Diagram Preferences on page 711](#) for a more detailed procedure.

### To Set Preferences:

1. From the Window menu, select **Preferences > TIBCO BusinessEvents**.
2. Expand the options on the left and select each option to view a panel of options on the right.  
Details on each set of options are provided in sections following.
3. Change the options from the list as desired, and then click **Apply**.
4. Click **OK** to see your preferred settings.

## Preference Sections

For tester and tester appearance preferences, see [Tester Preferences on page 735](#).

### TIBCO BusinessEvents (Top Level)

The top level preferences allow you do the following:

- Show or hide hover (tooltip) information in Catalog Functions view.
- Show or hide the confirm open perspective message pop-up.
- Switch or not switch to the default perspective every time an editor opens.
- Clear all the "do not show again" settings and show all hidden dialogs again.

### Build Path and Classpath Variables

Maintains a list of the classpath variables that have been defined. You can add, update, or remove classpath variables.

### Code Generation

- Compilation Mode options are File System or In Memory
- Source Java Version options are 1.5 or 1.6
- Target Java Version options are 1.5 or 1.6

## Code Generation and Ignored Resources

Maintain the list of ignored resource patterns here: Code Generation > Ignored Resources. These preferences define what file types are excluded from the EAR file shared archive section at build time.

You can add and remove patterns, and select which existing patterns to ignore or stop ignoring.

## Compare/Merge, Decision Table

Compare/Merge preferences are used for comparing two decision tables.

For Decision Table preferences see [Decision Table Related Preferences on page 728](#) for details.

Decision tables are available only when the TIBCO BusinessEvents Decision Manager add-on is in use..

## Diagram Preferences

In addition to a global preferences section, there are sections for each type of diagram. See [Diagram Preferences on page 730](#) for details.

## Error/Warning Preferences

You can enable and disable error and warning messages for each type of resource in a project. Resource extensions are shown so you can select which ones to enable or disable. By default all are checked (enabled). If you make changes, a prompt appears asking if you want to rebuild the project.

## Printing Preferences

These preferences apply to printing of diagrams, windows, and selections.

## Rules and Rule Debug Preferences

In addition to the standard Eclipse Run/Debug preferences, TIBCO BusinessEvents requires some rule debug preferences.

## Rules Management Server Preferences

These preferences are used with TIBCO BusinessEvents Decision Manager.

## Decision Table Related Preferences

---

### Decision Table Preferences

The Table View preferences set the initial state of the following features. Text in parentheses shows the names of corresponding buttons in the Decision Table editor. The settings can be toggled in the UI, but their initial state is set by preferences you set:

- Show column alias if present
- Automatically resize columns to fit content (Fit Content button)
- Automatically resize rows to fit content
- Automatically merge rows (Merge Rows button)
- Show expanded text (Show Text button)
- Show domain descriptions if present.
- Show column filter
- Show titled border
- Use existing IDs when importing
- Export column alias
- Automatically update decision tables on change in virtual rule function arguments
- Show condition area string
- The Use non-resizable editor sections option prevents or allows users from resizing editor sections.

### Decision Table Analyzer Preferences

You can control the following behavior of the decision table analyzer, as well as its appearance:

- Highlight partial ranges
- Use domain model for table completeness
- Show analyzer contents while opening table.

## Decision Table Appearance Preferences

These preferences let you set the color scheme for decision tables. For each item in the list, you can set color preferences. For condition data and action data you can specify a font.

### Compare/Merge Preferences

The tabbed dialog is in the TIBCO BusinessEvents preferences. It applies to decision tables.

In the Structure Compare tab (shown above), set preferences for initial settings and colors. The effect of your choices is shown in the lower panel.

In the Text Compare tab (shown below), you can set preferences for text. The effect of your choices is shown in the lower panel.

The Merge tab has one option: Allow columns to be automatically merged.

## Diagram Preferences

---

The diagram preferences apply to the following types of diagrams:

- Concept
- Dependency
- Event
- Process (This option is available only if TIBCO BusinessEvents Process Orchestration is installed.)
- Project
- Sequence
- State Model (This option is available only if TIBCO BusinessEvents Data Modeling is installed.)

Preferences set using a diagram's palette are applicable only to the displayed diagram.

Preferences set using the Preferences dialog define the default preferences for diagrams of that type.

Preferences for specific diagrams override global preferences, where the same settings exist.

Table 71 Global Diagram Preferences

| Option                   | Description                                                                                                                                                                               |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Reset Tool After Changes | If enabled, resets the tool to the Select tool after you add a node to the diagram.<br>Default: Enabled.                                                                                  |
| Auto Hide Scrollbars     | If enabled, then when the complete contents can display in the space available, scrollbars are hidden.<br>If not enabled, scrollbars remain.<br>Default: Auto hide scrollbars is enabled. |
| Show Tooltips            | If enabled, you can see the tooltips when you hover the mouse on an element in the diagram.<br>Default: Show Tooltips is enabled.                                                         |
| Link Types               | Shows the links as straight lines or curved lines.<br>Default: Straight.                                                                                                                  |

Table 71 Global Diagram Preferences (Cont'd)

| Option                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Run Layout On Changes | <p>After adding or deleting a node or an edge, refreshes the layout of the diagram. The layout options are None, Incremental, and Full.</p> <p>Default: None.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Animation             | <p>You can specify how a diagram behaves while it is changing from one layout option to another. Animation options allow you to see the nodes and edges moving from one arrangement to another. If you disable animation, the display simply switches from one kind of layout to another.</p> <p><b>Fade:</b> When enabled, the diagram transitions from faded out to fully colored state. The transition completes in the duration specified in the Layout duration option.</p> <p><b>Interpolation:</b> When enabled, the nodes and edges appear to move from one layout to another, when you select a different layout. The transition completes in the duration specified in the Layout duration option.</p> <p><b>Viewport Change:</b> Enables the interpolation animation when the viewport changes as a result of zooming. If enabled, the zoom change is reflected gradually with animation.</p> <p>By default, animation is enabled along with Fade, Interpolation and Viewport Change.</p> <p><b>Duration:</b> The time taken by animation to complete.</p> <p><b>Layout:</b> Time, in milliseconds, for the layout change to complete. Applicable to the interpolation and fade options.</p> <p>Default for Layout: 500</p> <p><b>Viewport Change:</b> Time, in milliseconds, for the viewport change to complete.</p> <p>Default for Viewport Change: 1000</p> |
| Opaque Movement       | <p>Interactive Zoom Sensitivity specifies the sensitivity of the mouse in interactive zooming.</p> <p>Pan Sensitivity specifies the sensitivity of the mouse in panning.</p> <p>Default: Opaque Movement is enabled, Interactive Zoom Sensitivity is 200.0, and Pan Sensitivity is 1.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Magnify Tool          | <p>The size the zoom window, and the level of zoom.</p> <p>Default: Window Size is 250, and Zoom Level is 3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

The preferences for all types of diagrams are listed in [Table 72 on page 732](#). The preferences are organized alphabetically and not grouped according to the diagram types.

Table 72 *Specific Diagram Preferences (Alphabetical) (Sheet 1 of 3)*

| Option                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Layout Style       | Sets whether to show Selected Entity Project diagrams in a circular or symmetric clustered layout.<br>Default: Circular.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Create view when analyzing | If enabled, creates a Selected Entity Project diagram when you choose to run the project analyzer.<br>Default: Disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Dependency Levels          | Sets how many levels of dependencies to view in a dependency diagram - One, two, or all.<br>Default: One.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Filter Options             | Sets whether to show various project resources in a Selected Entity Project diagram: Show Concepts, Show Events, Show Decision Tables, Show Domain Model, Show State Machines, Show Archives, Show Rules, Show Rule functions, Show Scorecard, Show Channels, Show Scope Links, Show Usage Links, Show Archived Destinations, Show Archived Rules, Show Archived Rules (All), Show Rules in Folders, Show Tooltips, Group Concepts, Group Events, Group Rules, and Group Rule Functions.<br>Default: All enabled except for Archived Rules, Group Concepts, Group Events, Group Rules, and Group Rule Functions. |
| Fix node and edge labels   | Keeps the labels of the nodes and edges with their graphics.<br>Default: Enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Grid                       | Sets the display either to without grid, or grid with lines, or grid with points.<br>Default: Lines.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Layout Style               | Sets either the orthogonal or hierarchical layout.<br>Default: Orthogonal.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 72 *Specific Diagram Preferences (Alphabetical) (Sheet 2 of 3)*

| Option                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Layout Quality                         | <p>Options are Draft (lowest quality), Medium, and Proof (highest quality). It takes longer to generate a higher quality diagram than a lower quality diagram. A different algorithm is used in each case.</p> <p>Default: Draft.</p>                                                                                                                                                                                                                                                                                                                                                                             |
| Link Routing                           | <p>Determines how links are routed when the hierarchical layout is used:</p> <p>Orthogonal— Routes the links using horizontal and vertical straight line segments (that is, using right angles)</p> <p>Polyline—Routes the links using straight line segments with arbitrary angles</p> <p>Overlapping lines are more likely with orthogonal routing than with polyline routing. With polyline routing the routing algorithm adds path nodes as needed to avoid overlapping lines.</p> <p>Note that the line segments can be joined using straight lines or curves in both cases.</p> <p>Default: Orthogonal.</p> |
| Link Routing—<br>Fix Node<br>Positions | <p>This setting affects link routing behavior for all layout options.</p> <p>If enabled, node positions do not change when you use the link routing feature.</p> <p>If disabled, link routing changes node positions as needed for clarity.</p> <p>Default: Disabled.</p>                                                                                                                                                                                                                                                                                                                                         |
| Link Routing—<br>Fix Node Sizes        | <p>This setting affects link routing behavior for all layout options except hierarchical layouts.</p> <p>If enabled, node sizes do not change when you use the link routing feature.</p> <p>If disabled, link routing changes node sizes as needed for clarity.</p> <p>Default: Disabled.</p>                                                                                                                                                                                                                                                                                                                     |
| Orientation                            | <p>Defines the general direction in which the links display, to reflect hierarchical relationships between the entities. TIBCO BusinessEvents diagrams are not particularly hierarchical, but setting this option defines the general direction of the layout.</p> <p>Options are: Top to Bottom, Bottom to Top, Left to Right, and Right to Left.</p> <p>Default: Top to Bottom.</p>                                                                                                                                                                                                                             |

Table 72 Specific Diagram Preferences (Alphabetical) (Sheet 3 of 3)

| Option                              | Description                                                                                                                                                                                                                                                     |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Orthogonal Fix Node Sizes           | This setting affects link routing behavior for the hierarchical layout option only.<br>If enabled, node sizes do not change when you use the link routing feature.<br>If disabled, link routing changes node sizes as needed for clarity.<br>Default: Disabled. |
| Run analysis when creating view     | Runs the Project Analyzer when creating a Selected Entity Project diagram.<br>Default: Enabled.                                                                                                                                                                 |
| Run fast layout for large diagrams  | Only for Selected Entity Project diagrams. If enabled, TIBCO BusinessEvents filters out certain properties before generating a large Selected Entity Project diagram to make it look simpler.<br>Default: Disabled.                                             |
| Show all properties in Concept Node | If enabled, a diagram shows all properties of a concept node, instead of showing only four default properties.<br>Default: Disabled.                                                                                                                            |
| Show all properties in Event Node   | If enabled, a diagram shows all properties of an event node, instead of showing only four default properties.<br>Default: Disabled.                                                                                                                             |
| Show Catalog Functions              | Sets whether to show catalog functions in a Sequence diagram.<br>Default: Enabled.                                                                                                                                                                              |
| Show Catalog Function Return Links  | Sets whether to show return links of catalog functions in a Sequence diagram.<br>Default: Disabled.                                                                                                                                                             |
| Show Expanded Names                 | Sets whether to show expanded names in a Sequence diagram.<br>Default: Enabled.                                                                                                                                                                                 |
| Snap to grid                        | If you move nodes in a diagram and the grid is shown, the nodes snap to the grid lines if Snap to grid is enabled.<br>Default: Disabled.                                                                                                                        |
| Undirected layout                   | Sets no orientation for the diagram.<br>Default: Disabled.                                                                                                                                                                                                      |

## Tester Preferences

---

The Tester preferences section lets you configure the tester behavior. See [Chapter 40, Testing and Debugging Projects, on page 679](#).

Table 73 Reference to Tester Preferences

| Option                             | Description                                                                                                                                                                               |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input Directory                    | The directory in which test data files you create are stored.<br>Default: /TestData                                                                                                       |
| Output Directory                   | The directory in which the result files are stored.<br>Default: /TestData/<Projectname>/<Processing Unit name>                                                                            |
| No of WM Objects                   | Maximum number of objects in working memory using a preference.<br>Default: 50                                                                                                            |
| Auto scroll Modified Result Tables | If enabled, you see scroll bars for the cells in Result Test Data after clicking Fit Content. These scroll bars move together for both the After and Before sections.<br>Default: Enabled |

When you test a project, tester shows the values changed while running the engine and the instances created in Result Test Data. To set the text color preferences for background and foreground of the modified values, go to **Window > Preferences > TIBCO BusinessEvents > Tester > Appearance**.

Table 74 Reference to Tester Appearance Preferences

| Option                           | Description                                                                                                                                        |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Background of the modified value | The background color of the changed values. Click the color box, which opens a color palette. Select the new color from it.<br>Default: Light Pink |
| Foreground of the modified value | The text color of the changed values. Click the color box, which opens a color palette. Select the new color from it.<br>Default: Blue             |

Table 74 Reference to Tester Appearance Preferences

| Option                      | Description (Cont'd)                                                                     |
|-----------------------------|------------------------------------------------------------------------------------------|
| Font for the modified value | The font in which the modified value appears. Click <b>Change...</b> to change the font. |
|                             | Default: Tahoma Regular 11                                                               |

---

## Appendix A Understanding Entity Caches

This appendix is provided for those who want to understand the internal structures of the Oracle Coherence caches used in Cache object management. This information is not required for configuration tasks.

For each entity in working memory, a corresponding cache exists in the cache cluster. Internal entities also have caches for various purposes, explained in this section.

### Entity Cache Names Format

Each entity cache has a name, which uses the following format:

*cache-type.cluster-name.AgentClassName.entity-name*

The elements of the above name are explained below

#### Cache Type (Caching Scheme)

Cache type is the type of caching scheme (as defined by its cache name in the `coherence-cache-config.xml` descriptor), for example, `dist-unlimited-bs`.

#### Cluster Name

Cluster name is the value of the following property:

```
java.property.tangosol.coherence.cluster
```

#### Agent Name

This field of the cache name is blank because TIBCO BusinessEvents does not support agent-specific entity caches.

All entities are globally scoped and available to all agents.

#### Entity Name

Two types of entities have caches:

- Internal entities
- Ontology entities

Internal entity names and caches are listed and described in [Table 75, Internal Entity Caches, on page 738](#).

The ontology entity field of the entity cache name uses the entity's generated class name, which is similar to its design-time folder path and name, prefixed by `be.gen`. For example:

```
be.gen.Concepts.LargeConcepts.ThisLargeConcept
```

## Caches for Ontology Objects

These caches are used to store the objects of types defined in the ontology of the project.

The types of caches created for ontology objects depend on the caching scheme used. If the `dist-unlimited-bs` caching scheme is used, then the cache names look like this:

```
dist-unlimited-bs$foo$$be.gen.Order
```

Where `foo` is the cluster name.

## Caches for Internal Entities

The following internal caches use a pre-defined scheme in the cache configuration file. Do not change this scheme. This information is provided for reference only.

Table 75 Internal Entity Caches

| Entity (Cache) name | Purpose of the Cache                                                                                                              |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| Master              | Maintains the cluster state and is shared by all nodes.                                                                           |
| Catalog             | Maintains a cached copy of all ontology definitions shared by all nodes.                                                          |
| TypeIDs             | TypeIDs are internal numbers used to store the short number to identify the URL of a concept instead the large string.            |
| ObjectTableIDs      | Stores the key mapping for all objects in the cluster. The objects themselves are stored in their respective caches.              |
| ObjectTableExtIDs   | Stores the external key mapping for all objects that have an external ID ( <code>extId</code> ).                                  |
| AgentTable          | Stores all the agents and their respective states across all cluster nodes and identifies the currently active and standby nodes. |

Table 75 Internal Entity Caches

| Entity (Cache) name     | Purpose of the Cache                                                                                                                                                                                                                                                                                     |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>AgentTxn-agentId</i> | Each agent in the cluster has an <i>AgentTxn-agentId</i> cache. The <i>agentId</i> is internally generated. It stores the change list for the agent. The change list is used to replicate changes between active-active and active-passive sets of agents in the cluster so that they stay synchronized. |
| TimeQueue               | Maintains all entries that are time bound, for example, state machines that can have timeouts at a state machine level or at a state level. This cache maintains an index to all objects that must be re-evaluated after a certain period of time.                                                       |



## Appendix B **Handling Null Properties**

This appendix explains three cases for special handling of null concept property values.

### Topics

---

- [Handling Null Concept Property Values, page 742](#)

## Handling Null Concept Property Values

---

This appendix explains three cases for special handling of null concept property values:

- [Enabling Use of the Nillable Attribute on page 742](#)
- [Enabling Null Property Values to Appear When Serializing Concepts to XML on page 742](#)
- [Enabling and Setting Special Treatment of Numeric and Boolean Null Values on page 744](#)

This section also provides related procedures and a reference to the properties used ([Table 76, Properties for Null Property Handling, on page 746](#)).

### Enabling Use of the Nillable Attribute

The presence of the `xsd:nillable` attribute in an XSD element means that the corresponding element in the XML file permits null values.

Setting `tibco.be.schema.nil.attrs=true` in `studio.tra` causes the `xsd:nillable` attribute ("`xsd:nillable=true`") to be set on all elements in the TIBCO BusinessEvents concept XSD. When an element in the XML file generated using that XSD has a null value, the `xsi:nil="true"` attribute is set on that element.

When set to false, the `xsd:nillable` attribute is not added and the corresponding XML file does not treat empty elements as null values.

In the absence of the `xsd:nillable` attribute in the XSD element, a corresponding empty element in the XML file is assumed to have a value. Elements that have no value are treated as empty strings ("").



**Effect on schema generation tool** The setting for this property affects the concept XSD files generated using the Generate Schema utility. See [Exporting \(Generating\) Concept and Event Schema \(XSD\) Files on page 5](#) for details.

### Enabling Null Property Values to Appear When Serializing Concepts to XML

By default concept properties with null values are excluded when concept objects (instances) are serialized to XML. You can override this behavior.

Setting the following property to false in the `studio.tra` file causes properties with null values to be included in the XML representation of a concept:

```
tibco.be.schema.exclude.null.props=false
```

## Examples of Nillable Attribute and Null Properties Settings

These examples illustrate the effect of the following properties on concept serialization:

```
tibco.be.schema.nil.attrs
tibco.be.schema.exclude.null.props
```

### If Null Properties are Excluded

```
tibco.be.schema.nil.attrs= true or false
tibco.be.schema.exclude.null.props=true
```

Suppose a Customer concept instance has no value for its CustomerName property. By default, the CustomerName property is excluded from the XML output. The output might look like the following:

---

```
<CustomerID>111</CustomerID>
<Country>Japan</Country>
<City>Tokyo</City>
```

---

If null properties are excluded when concepts are serialized, the `tibco.be.schema.nil.attrs` property has no effect on concept serialization.

### If Null Properties are Included and the Nillable Attribute is Set

```
tibco.be.schema.nil.attrs=true
tibco.be.schema.exclude.null.props=false
```

The output for the Customer concept instance shown above would be as follows, where there is no value for the CustomerName element in the concept instance:

---

```
<CustomerID>111</CustomerID>
<Country>Japan</Country>
<City>Tokyo</City>
<CustomerName
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:nil="true"/>
```

---

### If Null Properties are Included and the Nillable Attribute is not Set

```
tibco.be.schema.nil.attrs=false
tibco.be.schema.exclude.null.props=false
```

In this case, each null property is considered to be an empty string, and is represented, for example, as follows:

```
<CustomerName/>
```

## Enabling and Setting Special Treatment of Numeric and Boolean Null Values

If you enable null values to be output to XML (see [Enabling Null Property Values to Appear When Serializing Concepts to XML on page 742](#)), then you may also want to configure additional properties for defining how to treat null values for numeric types and Booleans, as explained in this section.

TIBCO BusinessEvents does not implicitly support null values for numeric types and Booleans. This can lead to interoperability issues when working with external sources, such as databases, which do permit blank (null) values.

To address such issues, you can enable special treatment of numeric null values at design time. At runtime, TIBCO BusinessEvents then uses a special numeric value for each numeric datatype, and (by default) `Boolean.FALSE` for Booleans, to represent a null value. Default special values are provided and you can override the defaults at runtime (see [Table 76, Properties for Null Property Handling, on page 746](#)).

The special numeric values that indicate null are used in TIBCO BusinessEvents when serializing and deserializing a concept to and from its XML representation, and when performing various operations on database concepts and the database tables to which they are linked.

The special values that indicate null appear only in TIBCO BusinessEvents. The appropriate null value is used in the XML or database representation of the concept property.

Conversely, when deserializing or importing a concept instance, TIBCO BusinessEvents represents null values in the source using these special values that indicate null in the concept instance.

To enable special treatment of numeric null values, set the following property in `studio.tra`:

```
tibco.be.schema.treat.null.values=true
```

At runtime you can override the default values that indicate null. See [Setting Runtime Properties for Special Treatment of Null Values on page 745](#).

### Summary

Set the following properties in `BE_HOME/studio/eclipse/configuration/studio.tra` as desired:

- To enable null property values to appear when serializing concepts to XML, add the following property and set it to false:  
`tibco.be.schema.exclude.null.props=false`
- To enable use of the nillable attribute in the concept XSD, add the following property and set it to true:  
`tibco.be.schema.nil.attrs=true`

- To enable special handling of null properties, add the following property and set it to true:

```
tibco.be.schema.treat.null.values=true
```

## Setting Runtime Properties for Special Treatment of Null Values

If you have enabled special treatment of null numeric properties, you can override the default special values that indicate null values of various datatypes in TIBCO BusinessEvents (shown in [Property Reference for Null Property Handling on page 746](#)).

To override the default values do the following:

1. In TIBCO BusinessEvents Studio, open the CDD editor for the project and add the following properties at the cluster level properties sheet. Provide the special values as desired:

```
tibco.be.property.int.null.value=value
```

```
tibco.be.property.long.null.value=value
```

```
tibco.be.property.double.null.value=value
```

```
tibco.be.property.boolean.null.value=value
```

Choose values that will not be misinterpreted as literal values.

2. Save the file.
3. Deploy the project with the updated CDD file.

## Property Reference for Null Property Handling

Set the following properties in the `studio.tra` file as needed to configure the output for your needs before you generate the EAR file.

Table 76 Properties for Null Property Handling

Property	Notes
<b>Properties Set in <code>BE_HOME/studio/eclipse/configuration/studio.tra</code></b>	
<code>tibco.be.schema.nil.attrs</code>	<p>Setting this property to true causes the <code>xsd:nilable</code> attribute ("<code>xsd:nilable=true</code>") to be set on all elements in the TIBCO BusinessEvents concept XSD.</p> <p>See <a href="#">Enabling Use of the Nillable Attribute on page 742</a></p> <p>Possible values are true and false.</p> <p>Default is false.</p>
<code>tibco.be.schema.exclude.null.props</code>	<p>When this property is set to true, null-valued concept properties are not output when the concept is serialized to XML.</p> <p>When set to false, null-valued concept properties are output to XML.</p> <p>See <a href="#">Enabling Null Property Values to Appear When Serializing Concepts to XML on page 742</a>.</p> <p>Possible values are true and false.</p> <p>Default is true.</p>
<code>tibco.be.schema.treat.null.values</code>	<p>Setting this property to true causes TIBCO BusinessEvents to use special numeric values that indicate null for numeric datatypes. The special numeric values are set using the properties listed next.</p> <p>See <a href="#">Enabling and Setting Special Treatment of Numeric and Boolean Null Values on page 744</a>.</p> <p>Possible values are true and false.</p> <p>Default is false.</p>

Table 76 Properties for Null Property Handling (Cont'd)

Property	Notes
<b>Properties Set in CDD at the cluster level</b>	
<pre>tibco.be.property.int.null.value tibco.be.property.long.null.value tibco.be.property.double.null.value tibco.be.property.boolean.null.value</pre>	<p data-bbox="392 430 1316 489">These properties define a special value that indicates null. Use a value that will not be confused with an actual numeric value.</p> <p data-bbox="392 512 1316 571">These properties are used only if <code>tibco.be.schema.treat.null.values</code> is set to true.</p> <p data-bbox="392 593 1266 621">Default values for each numeric datatype are the following Java constants:</p> <pre data-bbox="392 644 711 671">int: Integer.MIN_VALUE</pre> <pre data-bbox="392 692 682 720">long: Long.MIN_VALUE</pre> <pre data-bbox="392 741 739 769">double: Double.MIN_VALUE</pre> <pre data-bbox="392 789 711 817">boolean: Boolean.FALSE</pre> <p data-bbox="392 822 1285 881">For Integer and Long these constants represent the most negative value. For Double the constant represents smallest positive nonzero value (4.9e-324).</p>



# Index

## Symbols

[@category](#) 190  
[@closure](#) 182, 329  
[@extId](#) 173, 190, 200, 329  
[@id](#) 173, 182, 190, 200, 329  
[@interval](#) 182, 329, 329  
[@isSet](#) 330  
[@isset](#) 201  
[@length](#) 330  
[@message](#) 190  
[@parent](#) 200, 330  
[@payload](#) 173, 329  
[@scheduledTime](#) 182, 329, 329  
[@ttl](#) 173, 183, 329, 329, 329  
[@type](#) 190

## A

[Accept Count](#) 108  
[acknowledgement modes](#) 89  
[action-only functions](#) 302  
[actions](#) 253
 

- [in rule editor](#) 259, 288

[activities](#)

- [Input tab icons](#) 349

[activities palette](#) 638, 645  
[Adding a Cluster Deployment Descriptor](#) 403  
[Adding a JDBC Connection Resource to the Studio Project](#) 537  
[Adding Business Rules](#) 617  
[addressing schema elements](#) 384  
[advisory events](#) 188, 188  
[AdvisoryEvent event type](#) 334  
[agent key, and BEMM](#) 516  
[Agent.AgentClassName.cacheTxn.updateCache](#) 484, 502

[Agent.AgentClassName.checkDuplicates](#) 503  
[Agent.AgentClassName.dbOpsBatchSize](#) 465, 504  
[Agent.AgentClassName.dbOpsQueueSize](#) 504  
[Agent.AgentClassName.dbthreadcount](#) 460, 504  
[Agent.AgentClassName.enableParallelOps](#) 495, 502, 503  
[Agent.AgentClassName.recoveryPageSize](#) 501  
[Agent.AgentClassName.threadcount](#) 502, 502  
[Agent.AgentGroupName.checkDuplicates](#) 173, 200  
[Agents Tab Settings and Properties](#) 498  
[aliases](#) 259  
[all values history policy](#) 199  
[Approval Status Values](#) 626  
[Approving and Rejecting Commit Requests](#) 626  
[array indexes, start from zero or one](#) 268, 332  
[arrays](#)

- [accessing and appending values](#) 332
- [indexes start from zero or one](#) 268, 332
- [primitive](#) 324

[attributes](#) 329  
[AUTO\\_ACKNOWLEDGE](#) 89

## B

[backing store](#) 397, 523
 

- [cache aside](#) 458
- [configuring](#) 523
- [database type](#) 457
- [excluding entities from](#) 487
- [Initial Size \(backing store connection pool\)](#) 460
- [Max Size \(backing store connection pool\)](#) 460
- [Min Size \(backing store connection pool\)](#) 460
- [strategy \(for Oracle DB in JDBC backing store\)](#) 457
- [updating schema of](#) 552
- [URI](#) 459

[Backing Store Table Reference](#) 556  
[BasicMessageSerializer](#) 86

BE\_HOME xxxvii  
 be.agent.query.enable.filter.optimizer 505  
 be.agent.query.localcache.prefetchaggressive 505  
 be.auth.file.location 577  
 be.auth.ldap.usessl 580  
 be.auth.type 577  
 be.backingstore.commitSize 463  
 be.backingstore.dburi.pool.inactivityTimeout.0 462  
 be.backingstore.dburi.pool.waitTimeout.0 461  
 be.backingstore.readtimeout.0 462  
 be.backingstore.recreateOnRecovery 462  
 be.channel.deactivate 572  
 be.channel.tibjms.queue.ack.mode 89  
 be.channel.tibjms.queue.disabled 521  
 be.channel.tibjms.topic.ack.mode 89  
 be.channel.tibjms.topic.disabled 521  
 be.cluster.ruletemplateinstances.deploy.dir 636  
 be.codegen.rootDirectory 575  
 be.codegen.useLegacyCompilation 575, 576  
 be.engine.cacheServer.channel.disable 505  
 be.engine.cluster.as.discover.url 419, 421, 430  
 be.engine.cluster.as.listen.url 431  
 be.engine.cluster.recovery.batchsize 535  
 be.engine.cluster.recovery.threads 456, 479  
 be.engine.cluster.useDBBatching 465  
 be.engine.profile.\*.duration 670  
 be.engine.profile.\*.enable 668  
 be.engine.profile.\*.file 669  
 be.engine.profile.\*.level 670  
 be.engine.profile.BAR\_Name.duration 669  
 be.engine.profile.BAR\_Name.enable 668  
 be.engine.profile.BAR\_Name.file 669  
 be.engine.profile.BAR\_Name.level 670  
 be.engine.profile.delimiter 667  
 be.jms.reconnect.msgCodes 521  
 be.jms.reconnect.timeout 520  
 be.mm.cluster.as.listen.url 432  
 BEMM, and agent key 516  
 BEReceiveEventOutput 647, 651  
 BESendEventInput 648  
 buildEar command line tool 26  
 buildLibrary command line tool 31  
 Business rule builder  
     command builder 616  
     condition builder 616

Business Rule Hot Deployment 636  
 BusinessWorks Repo URL 500  
 BytesMessageSerializer 86

## C

Cache Agent Quorum 411  
 cache aside 458  
 cache clusters  
     localport auto incrementing feature 440, 441, 443, 443  
 cache provider 411  
 Caches for Internal Entities 738  
 Caches for Ontology Objects 738  
 caller (threading model) 495  
     and order of acknowledgements 495  
 Candidate Event Key 650  
 CDD  
     adding 403  
     editing after deployment 397  
     how settings apply at runtime 400  
     working in the editor 396  
 CDD, setting global variables in 402  
 changes from the previous release of TIBCO BusinessEvents Developer's Guide xxvi  
 changes only history policy 199  
 ChannelURI, variable 88  
 Check for Duplicates 500  
 Checking a Worklist and Taking Action 628  
 Checking Out a Project 588  
 circular layout 720, 722  
 Class-to-Table Mapping 558  
 CLIENT\_ACKNOWLEDGE 89, 91  
 ClientID (channel property) 75  
 cluster provider setting 411  
 Cluster Tab — Cache OM — Backing Store  
     Properties 461  
 Cluster Tab — Cache OM — Backing Store  
     Settings 457  
 Cluster Tab — Cache OM — Coherence Properties 438  
 Cluster Tab — Cache OM — Configuration  
     Settings 411  
 Cluster Tab — Cache OM — Domain Objects

- Settings 481, 485
- Cluster Tab — General Settings 405
- cmigrating core Coherence functions in TIBCO BusinessEvents Studio 50
- code generation, ignored resources 727
- Coherence
  - enabling as cache provider 433
  - enabling functions 433
- Coherence functions
  - migrating at command line 34
  - migrating in TIBCO BusinessEvents Studio 50
- coherence-cache-config-jdbc.xml 434
- Collections Tab — Destinations Settings and Properties 494
- Collections Tab — Log Configurations Overview 506
- Collections Tab — Log Configurations Settings 509
- com.tibco.cep.runtime.channel.payload.validation 501
- com.tibco.tibjms.connect.attempts 520
- Committing Artifacts for Approval 591
- complex element 170
- Compressible Mime Types 113
- Compression 110
- concept properties, accessing 331
- concepts
  - and ontology functions 299
  - overview 194
- concurrent RTC 500, 503
- concurrentwm (concurrent RTC) 500
- conditions 253
  - in rule editor 259
- Conditions and Actions 600
- Configuring Basic RMS Server Properties 569
- Configuring for Deployment of Decision Table Classes 632
- Configuring Log Configurations 507
- Configuring RMS Server Properties 569
- Configuring the Agent Classes Tab (All Object Managers) 496
- Configuring the Collections Tab (All Object Managers) 492
- Configuring the Processing Units Tab (All Object Managers) 513
- Connection Linger 109
- Connection Timeout 108

- constants in XPath 355
- controlling start of state machines 198
- Create/Remove a Decision Table 596
- Creating a Rule Management Server Project
  - Update a Project 590
- custom functions
  - name overloading not supported 296
  - return types supported 296
  - static modifiers 296
- Custom ID 647
- customer support xl

**D**

- daemon (channel property) 74
- Database Concepts 404
- Database Connection Settings 552
- Database Location 531
- database type, for Shared All Persistence 457
- date functions 291
- DateTime functions 291
- dbkeywordmap.xml (for JDBC backing store) 529
- DBMS Software and Installation Requirements 531
- Debugger
  - global variables in 18
- Decision Manager and RMS Basic Workflow 566
- declaration 252
  - rule editor 259, 287
- default events 61
- Default Level Settings 481, 485
- Deleting Business Rules 620
- Deployment category of advisory event 191
- design-time log file 679
- destination queue (threading model) 495
- destinations
  - and channels 58, 78
  - default destination for event 168
- DestinationURI, variable 88
- diagrams
  - layout menu options 720
  - Link Navigator 720
  - preferences 730
- Document Page 113, 114

Document Root 113  
 domain objects 397  
 DUPS\_OK\_ACKNOWLEDGE 90  
 durable subscriber name, and special variables 87  
 DurableSubscriberName 87

## E

EAR files, building at command line 26  
 Edit a Decision Table 597  
 Email Notification Property Group 581  
 Enable Cache Storage 515  
 Enable DNS Lookups 109  
 Enabling Remote Connection to RMS 571  
 Engine (advisory) 190  
 Engine category of advisory event 191  
 engine functions 292  
 engine properties  
   See also Index of Engine Properties  
 engine.primary.activated (advisory) 191  
 Engine.Profiler.startCollectingToFile() 671  
 Engine.Profiler.stopCollecting() 671  
 EngineName, variable 87  
 engines  
   configuring for remote debugging 683  
 Entity Cache Names Format 737  
 Entity Cache Size 412  
 Entity Level Object Overrides and Settings 485  
 ENV\_HOME xxxvii  
 errors 348, 348  
 errors in mappings 348  
 escape sequences 326  
 evaluation context 385  
 event functions 292  
 event properties, accessing 333  
 event scheduler functions 185  
 Event.routeTo 168  
 Event.sendEvent 168

events 168  
   advisory 188, 188  
   and ontology functions 299  
   default 61  
   default destination 168  
   inherits from, setting 167  
   naming restrictions 169, 180  
   payload parameters 170  
   scheduling 185  
 Eviction Time (Local Cache) 498  
 Example Alter Script 554  
 Exception (advisory) 190  
 Exception category of advisory events 190  
 Exception function 292  
 Exception Tables 563  
 excluding entities from the backing store 487  
 Expiry Action 169  
 explicit acknowledgement 90, 90  
 EXPLICIT\_CLIENT\_ACKNOWLEDGE 90  
 EXPLICIT\_CLIENT\_DUPS\_OK\_ACKNOWLEDGE 90  
 0  
 Export a Business Rule 620  
 Export a Decision Table 598  
 extended functions 300  
 extId  
   must be unique 173

## F

Fault Tolerance 106  
 File functions 292  
 formats for filter values 154  
 functions  
   action only 302  
   extended (hidden) 300  
   mapper 266, 302  
   standard 291  
   temporal 303  
   tool tips for 301  
   types and usage 290  
 functions documentation, accessing xxxvi

**G**

generateClass command line tool [35](#)  
 Generating Deployable Files [632](#)  
 global variables  
   in Debugger [18](#)  
   in the rule editor [18, 265](#)  
   setting in the CDD [402](#)  
   setting in TIBCO BusinessEvents Studio [15](#)  
   using in fields and rule editor [18](#)

**H**

Has Backing Store [487](#)  
 hidden functions [300](#)  
 hierarchical layout [723](#)  
 hierarchical layout with normal routing [721](#)  
 hierarchical layout with orthogonal routing [720](#)  
 hints [347](#)  
 History [199](#)  
 HTTP functions [292](#)

**I**

identifier requirements [323](#)  
 ignored resources [727](#)  
 Import a Decision Table [599](#)  
 importing a 4.x project at command line [28](#)  
 incremental layout [721, 723](#)  
 Inherits From [197](#)  
 inherits from event setting [167](#)  
 Initial Size (backing store connection pool) [460](#)  
 initialize\_database.sql [543](#)  
 initialize\_database.sql (for JDBC backing store) [530](#)  
 Input tab  
   icons [349](#)  
 Insert Model Group Content [358, 358, 358](#)  
 Instance.isModified() function, and scorecards [204](#)  
 Instance.startStateMachine() [198](#)  
 instances  
   aliases for [259](#)

Introduction to Decision Manager [560](#)  
 Introduction to Rules Management Server (RMS) [560, 568](#)  
 INVOKE BW PROCESS type of advisory event [190](#)  
 INVOKE BW PROCESS, type of BEBW advisory event [659](#)  
 IO functions [292](#)  
 isTransacted (channel property) [75](#)

**J**

Java Debug Interface (JDI) for remote debugging [683](#)  
 java.security.auth.login.config [577](#)  
 javax.net.ssl.trustStore [580](#)  
 javax.net.ssl.trustStorePassword [580](#)  
 javax.net.ssl.trustStoreType [581](#)  
 JDBC Backing Store Database Configuration Tasks [543](#)  
 JDBC Backing Store Database Setup Overview [524](#)  
 JMS  
   channel naming restrictions [169, 180](#)  
   converting messages to non-default events [61](#)  
   durable subscriber name [87](#)  
   header properties [169, 180](#)  
   message acknowledgement modes [89](#)  
   message acknowledgment [92](#)  
   reconnect attempts [520](#)  
   reconnect codes [521](#)  
   reconnect timeout [520](#)  
 JMSCorrelationID [96](#)  
 JMSDeliveryMode [95](#)  
 JMSDestination [95](#)  
 JMSExpiration [95](#)  
 JMSMessageID [96](#)  
 JMSPriority [95](#)  
 JMSRedelivered [96](#)  
 JMSReplyTo [96](#)  
 JMSTimestamp [96](#)  
 JMSType [96](#)

**K**

Key word mapping file 549

**L**

L1Cache 498

layout

labeling 721

layout menu 720

layout options

circular 722

hierarchical 723

incremental 723

orthogonal 722

symmetric 722

Link Navigator 720

Link Routing 721

links in diagrams, view path of 720

local channels

maximum events in queue 81

local variables 324

localhost 440

localport 440

auto incrementing feature 440, 441, 443, 443

lockWM 657

log configurations

adding 507

for TIBCO DataGrid 507, 509

levels 506

understanding 506

log file, design-time 679

log files

debug log (system functions for) 292

location of (setting) 506, 509

log properties

for Oracle Coherence 517

**M**

mapper functions 266, 302

mapping

addressing schema elements 384

Input tab icons 349

XPath operators and functions 389

mappings 348

math functions 292

Max Active 499

Max HTTP Header Size 111

Max HTTP Post Size 111

Max HTTP Save Post Size 112

Max KeepAlive Requests 111

Max Processors 109

Max Size (backing store connection pool) 460

Max Size (Local Cache) 498

Max Spare Threads 112

maximum (history parameter) 199

Maximum Connections, JDBC resource 460, 537

migrateCoherenceCalls 34

migrating core Coherence functions at command line 34

Min Size (backing store connection pool) 460

Min Spare Threads 112

Minimum User Permissions 531

Multiple (array) 199

**N**

Names that Exceed the DBMS Maximum Column Length 527

namespace 61

naming requirements 323

NCName datatype 127, 403

network (channel property) 74

New property 554

New table 554

NO\_ACKNOWLEDGE 91

number functions 292

Number of Backup Copies 412

## O

- Object Table Cache Size [412](#)
- One-way SSL Authentication in RMS [572](#)
- ontology functions
  - location of [299](#)
  - purpose and types of [299](#)
- Ontology Identifiers That Use Database Key Words [528](#)
- operators supported in filters [152](#)
- Oracle [534](#)
- ORACLE cache provider, setting [411](#)
- Orthogonal Layout [720](#)
- orthogonal layout [722](#)

## P

- palettes
  - TIBCO BusinessEvents activities [638](#), [645](#)
- Parallel operations property [495](#), [502](#), [503](#)
- payload parameters
  - cardinality [170](#)
  - complex element [170](#)
- payloads [170](#)
  - all, parameter [171](#)
  - attribute of type, parameter [171](#)
  - choice, parameter [171](#)
  - element of type, parameter [170](#)
  - sequence, parameter [171](#)
  - validation, parameter [172](#)
  - XML element reference, parameter [170](#)
  - XML group reference, parameter [171](#)
- PGM (Pragmatic General Multicast) URL Format [419](#)
- Policy (concept history) [199](#)
- port conflicts, avoiding [440](#)
- pre-conditions
  - in rule template editor [287](#)
- preferences
  - ignored resources [727](#)
- preprocessors
  - options on failure [168](#)
- Primary Tables [556](#)
- primitive arrays [324](#)

- Priority [258](#)
- Processing Units Tab Settings and Properties [515](#)
- processTimeout [650](#)
- project libraries
  - working with at command line [31](#)
- properties
  - Agent.AgentClassName.cacheTxn.updateCache* [484](#), [502](#)
  - Agent.AgentClassName.checkDuplicates* [503](#)
  - Agent.AgentClassName.dbOpsBatchSize* [465](#), [504](#)
  - Agent.AgentClassName.dbOpsQueueSize* [504](#)
  - Agent.AgentClassName.dbthreadcount* [460](#), [504](#)
  - Agent.AgentClassName.enableParallelOps* [495](#), [502](#),

503

Agent.*AgentClassName*.recoveryPageSize 501  
 Agent.*AgentClassName*.threadcount 502  
 be.agent.query.enable.filter.optimizer 505  
 be.agent.query.localcache.prefetchaggressive 505  
 be.backingstore.commitSize 463  
 be.backingstore.dburi.pool.inactivityTimeout.0 462  
 be.backingstore.dburi.pool.waitTimeout.0 461  
 be.backingstore.optimize.reads 463  
 be.backingstore.optimize.writes 463  
 be.backingstore.readtimeout.0 462  
 be.backingstore.recreateOnRecovery 462  
 be.backingstore.timestamp.useDateTimeZone 463  
 be.channel.tibjms.queue.disabled 521  
 be.channel.tibjms.topic.disabled 521  
 be.engine.cacheServer.channel.disable 505  
 be.engine.cluster.as.discover.url 419, 421, 430  
 be.engine.cluster.as.listen.url 431  
 be.engine.cluster.cleanup 464  
 be.engine.cluster.recovery.threads 456, 464, 479  
 be.engine.cluster.useDBBatching 465  
 be.engine.profile.\*.duration 670  
 be.engine.profile.\*.enable 668  
 be.engine.profile.\*.file 669  
 be.engine.profile.\*.level 670  
 be.engine.profile.*AgentClassName*.duration 669  
 be.engine.profile.*AgentClassName*.enable 668  
 be.engine.profile.*AgentClassName*.file 669  
 be.engine.profile.*AgentClassName*.level 670  
 be.engine.profile.delimiter 667  
 be.jdbc.schemamigration.pswd 552  
 be.jdbc.schemamigration.url 552  
 be.jdbc.schemamigration.user 552  
 be.jms.reconnect.msgCodes 521  
 be.jms.reconnect.timeout 520  
 com.tibco.cep.runtime.channel.payload.validation

501

com.tibco.tibjms.connect.attempts 520  
 tangosol.coherence.clusteraddress 438  
 tangosol.coherence.clusterport 438  
 tangosol.coherence.distributed.threads 442  
 tangosol.coherence.localhost 439  
 tangosol.coherence.localport 440  
 tangosol.coherence.log 517  
 tangosol.coherence.log.level 518  
 tangosol.coherence.log.limit 518  
 tangosol.coherence.management 519  
 tangosol.coherence.management.remote 519  
 tangosol.coherence.ttl 439  
 tangosol.coherence.wkan 435, 441  
 tangosol.coherence.wkan.port 435, 441  
 tibco.be.property.boolean.null.value 747  
 tibco.be.property.double.null.value 747  
 tibco.be.property.int.null.value 747  
 tibco.be.property.long.null.value 747  
 tibco.be.schema.exclude.null.props 746  
 tibco.be.schema.nil.attrs 746  
 tibco.be.schema.treat.null.values 746  
 property arrays, index from zero or one 268, 332  
 Property type change 554  
 property values, accessing 331, 331  
 Protocol (for HTTP channel) 112  
 Provided Configuration Resources 529  
 Provider setting for Cache OM 411  
 ProviderURL (channel property) 75

## Q

queue  
   local channel 81  
 Queue Size (Shared Queue) 499

## R

reconnect to JMS, attempts 520  
 reconnect to JMS, which messages 521  
 Regular and Custom Conditions and Actions 600

- remote debugging
  - configuring engine for 683
  - Java Debug Interface and 683
- Resources Required for Setting Up the Database 529
- RESTMessageSerializer 104
- Restricted User Agents 113
- Rete network 167, 188, 275, 309, 397
  - concurrent RTC feature 500, 503
- Retry On Exception 168
- Reverse Reference Tables 557
- Reverse References 488
- revert to an earlier version of a resource 44
- RMS Server Configuration Property Reference 575
- rms.checkin.revisionId.initValue 575
- rms.external.entities.autodetect 577
- rms.lockworkflowstages.config.file 576
- rms.project.decisiondata 576
- rms.project.deployment 576
- rms.project.workspace 576
- rms.roleArtifactTypes.config.file 577
- rms.workflowstages.config.file 576
- RTC 397
  - concurrent RTC 500, 503
- Rule Building with Decision Tables 562
- Rule Debugger
  - configuring engine for remote debugging 683
- rule editor
  - alias 259
  - global variables in 18, 265
- rule functions
  - ontology functions 299
- rule template 616
- rule templates
  - pre-conditions 287
- rules
  - actions 259, 288
  - and advisory events 188
  - and ontology functions 299
  - conditions 259
  - declaration 259, 287

## S

- scheduling events 185
- schema elements 346
  - addressing 384
- scorecards, and Instance.isModified() 204
- search predicates in XPath expressions 385
- Secondary Tables 557
- sending to another application 168
- Sequencing Key 647
- serializers 60
- service (channel property) 74
- Set Decision Table Properties 597
- shared queue 495
  - queue size 499
  - thread count 499
- SOAP functions 292
- Socket Output Buffer Size 108
- source control systems 3
- Special Cases 527
- specifying constants in XPath expressions 355
- SQL Server 532
- SSL Authentication in RMS 572
- standard functions 291
  - categories of 290
- StartFileBasedProfiler() 672
- Starting and Logging in to RMS 586
- state machines
  - controlling start of 198
- Status Check for RMS Connection 586
- StopFileBasedProfiler() 672
- Strategy (backing store) 457
- string functions 292
- String Properties That Exceed the Maximum Database Column Length 528
- studio-tools
  - core buildEar 26
  - core buildLibrary 31
  - core generateClass 35
  - core importExistingProject 28
  - core migrateCoherenceCalls 34
- subscription preprocessor required signature 487
- support, contacting xl
- symmetric layout 720, 722
- system functions 292

## T

Table Analyzer 563

tangosol.coherence.clusteraddress 438

tangosol.coherence.clusterport 438

tangosol.coherence.distributed.threads 460, 460

tangosol.coherence.guard.timeout 443

tangosol.coherence.localhost 439

tangosol.coherence.localport 440

tangosol.coherence.localport.adjust 440, 441, 443

tangosol.coherence.log 517

tangosol.coherence.log.level 518

tangosol.coherence.log.limit 518

tangosol.coherence.management 519

tangosol.coherence.management.remote 519

tangosol.coherence.ttl 439

tangosol.coherence.wkan 435, 441

tangosol.coherence.wkan.port 435, 441

tangosol-coherence-override-tibco-be.xml 434

TCP No Delay 109

technical support xl

temporal functions 293

- arguments 303

test connection feature, enabling 219

TextMessageSerializer 86

The Procedure 552

Thread Count (Shared Queue) 499

threading models 495

TIBCO BusinessEvents Virtual Rule Functions and Decision Tables 561

TIBCO BusinessWorks

- process definitions, and TIBCO BusinessEvents activities 638, 645

TIBCO cache provider, setting 411

TIBCO Rendezvous

- converting messages to non-default events 61

TIBCO\_HOME xxxvii

tibco.be.property.boolean.null.value 747

tibco.be.property.double.null.value 747

tibco.be.property.int.null.value 747

tibco.be.property.long.null.value 747

tibco.be.schema.exclude.null.props 746, 746

tibco.be.schema.nil.attrs 746

tibco.be.schema.treat.null.values 746

tibco.clientVar.RMS/Approval/adminRole 578

tibco.clientVar.RMS/hostname 578

tibco.clientVar.RMS/port 571, 578

tibco.clientVar.RMS/security/securePort 580

tibco.clientVar.RMS/security/sslCertificateStore 580

tibco.clientVar.RMS/security/sslCertificateStorePassword 580

tibco.clientVar.Webstudio/Connection/timeout 579

tibco.clientVar.Webstudio/DB/driverName 582

tibco.clientVar.Webstudio/DB/maxConnections 582

tibco.clientVar.Webstudio/DB/password 582

tibco.clientVar.Webstudio/DB/timeout 583

tibco.clientVar.Webstudio/DB/url 582

tibco.clientVar.Webstudio/DB/userName 582

tibco.clientVar.Webstudio/hostname 579

tibco.clientVar.Webstudio/port 579

tibco.clientVar.Webstudio/sessionTimeout 579

tibco.clientVar.Webstudio/warDir 579

time events

- and ontology functions 299

Time to Live 168

timeouts

- JMS reconnect 520
- local channels 81

tool tips

- creating for custom functions 295
- turning off display 301

tool tips, for functions 301

Two-way SSL Authentication in RMS 573

## U

Understanding Entity Caches 737

Updating (Refreshing) a Project 590

Updating an Existing Backing Store Database Schema 552

URI (for backing store) 459

URI Encoding 111

Use Body Encoding for URI 110

Using Agent Class Properties at Different Levels 400

UtfBytesMessageSerializer 86

Util functions 293

**V**

validation of XML payload elements [172](#)  
 variables for use with durable subscriber name [87](#)  
 variables, local [324](#)  
 VRF (Virtual Rule Function) functions [293](#)

**W**

What the Schema Update Utility Can and Cannot  
 Handle Automatically [553](#)  
 Workflow [566](#)  
 Working With Decision Tables [596](#)  
 ws.artifact.deploy.location [575](#)  
 ws.notify.enabled [581](#)  
 ws.notify.mail.domain [581](#)  
 ws.notify.prop.MAIL\_PROTOCOL [581](#)  
 ws.notify.prop.MAIL\_SERVER\_HOST [581](#)  
 ws.notify.prop.MAIL\_SERVER\_PORT [581](#)  
 ws.notify.prop.SENDER\_EMAIL [582](#)  
 ws.notify.prop.SENDER\_PASSWORD [582](#)  
 ws.notify.prop.SENDER\_USERNAME [582](#)  
 ws.projects.acl.location [568, 578](#)  
 ws.projects.customLib.location [576](#)  
 ws.scs.command.path [583](#)  
 ws.scs.default.password [584](#)  
 ws.scs.default.username [584](#)  
 ws.scs.impl.type [583](#)  
 ws.scs.repo.dir [583](#)  
 ws.scs.rootURL [568, 583, 584](#)  
 ws.validateDT.temp.dir [579](#)  
 WSDL  
 how web service clients view [122](#)

**X**

XML payloads element validation [172](#)

XPath [383](#)  
 basics [384](#)  
 editor [387](#)  
 evaluation context [385](#)  
 example [389](#)  
 operators and functions [389](#)  
 search predicates [385](#)  
 specifying constants [355](#)  
 XPath functions [293](#)  
 XSLT statements [347, 348](#)  
 XSLT template [266, 302](#)