



# **TIBCO EBX®**

## **Product Documentation**

*Version 5.9.26*  
*May 2024*



# Table of contents

---

## User Guide

---

### Introduction

|                                  |    |
|----------------------------------|----|
| 1. How TIBCO EBX® works.....     | 11 |
| 2. Using the user interface..... | 15 |
| 3. Glossary.....                 | 21 |

### Data models

|                                     |    |
|-------------------------------------|----|
| 4. Introduction to data models..... | 32 |
|-------------------------------------|----|

#### *Implementing data models*

|                                               |    |
|-----------------------------------------------|----|
| 5. Creating a data model.....                 | 37 |
| 6. Configuring the data model.....            | 39 |
| 7. Implementing the data model structure..... | 45 |
| 8. Properties of data model elements.....     | 51 |
| 9. Data validation controls on elements.....  | 65 |
| 10. Toolbars.....                             | 73 |
| 11. Working with an existing data model.....  | 81 |

#### *Publishing and versioning data models*

|                                            |    |
|--------------------------------------------|----|
| 12. Publishing a data model.....           | 83 |
| 13. Versioning an embedded data model..... | 85 |

### Dataspaces

|                                          |     |
|------------------------------------------|-----|
| 14. Introduction to dataspace.....       | 88  |
| 15. Creating a dataspace.....            | 91  |
| 16. Working with existing dataspace..... | 93  |
| 17. Snapshots.....                       | 101 |

### Datasets

|                                         |     |
|-----------------------------------------|-----|
| 18. Introduction to datasets.....       | 106 |
| 19. Creating a dataset.....             | 109 |
| 20. Viewing table data.....             | 111 |
| 21. Editing data.....                   | 119 |
| 22. Working with existing datasets..... | 121 |
| 23. Dataset inheritance.....            | 125 |

### Workflow models

|                                                     |     |
|-----------------------------------------------------|-----|
| 24. Introduction to workflow models.....            | 130 |
| 25. Creating and implementing a workflow model..... | 135 |
| 26. Configuring the workflow model.....             | 147 |
| 27. Publishing workflow models.....                 | 155 |

### Data workflows

|                                                       |     |
|-------------------------------------------------------|-----|
| 28. Introduction to data workflows.....               | 158 |
| 29. Using the Data Workflows area user interface..... | 159 |
| 30. Work items.....                                   | 165 |

### ***Managing data workflows***

|                                                  |     |
|--------------------------------------------------|-----|
| 31. Launching and monitoring data workflows..... | 171 |
| 32. Administration of data workflows.....        | 173 |

### **Data services**

|                                        |     |
|----------------------------------------|-----|
| 33. Introduction to data services..... | 178 |
| 34. Generating data service WSDLs..... | 181 |

# Reference Manual

---

## Integration

|                                                |     |
|------------------------------------------------|-----|
| 35. Overview of integration and extension..... | 187 |
| 36. Using TIBCO EBX® as a Web Component.....   | 191 |
| 37. Built-in user services.....                | 199 |

### *File import and export services*

|                                 |     |
|---------------------------------|-----|
| 38. XML import and export.....  | 213 |
| 39. CSV import and export.....  | 219 |
| 40. Supported XPath syntax..... | 225 |

## Localization

|                                                    |     |
|----------------------------------------------------|-----|
| 41. Labeling and localization.....                 | 234 |
| 42. Extending TIBCO EBX® internationalization..... | 237 |

## Persistence

|                                  |     |
|----------------------------------|-----|
| 43. Overview of persistence..... | 240 |
| 44. Relational mode.....         | 243 |
| 45. History.....                 | 249 |
| 46. Replication.....             | 257 |
| 47. Data model evolutions.....   | 263 |

## Other

|                                           |     |
|-------------------------------------------|-----|
| 48. Inheritance and value resolution..... | 268 |
| 49. Permissions.....                      | 273 |
| 50. Criteria editor.....                  | 291 |
| 51. Performance guidelines.....           | 293 |

# Administration Guide

---

|                                  |     |
|----------------------------------|-----|
| 52. Administration overview..... | 304 |
|----------------------------------|-----|

## Installation & configuration

|                                 |     |
|---------------------------------|-----|
| 53. Supported environments..... | 308 |
| 54. Java EE deployment.....     | 315 |

### *Installation notes*

|                                                                          |     |
|--------------------------------------------------------------------------|-----|
| 55. Installation note for JBoss EAP 7.4.x.....                           | 325 |
| 56. Installation note for Tomcat 9.x.....                                | 331 |
| 57. Installation note for WebSphere Application Server Liberty 23.x..... | 337 |
| 58. Installation note for WebLogic 14c.....                              | 343 |
| 59. TIBCO EBX® main configuration file.....                              | 349 |
| 60. Initialization and first-launch assistant.....                       | 369 |
| 61. Deploying and registering TIBCO EBX® add-ons.....                    | 371 |

## Technical administration

|                                          |     |
|------------------------------------------|-----|
| 62. Repository administration.....       | 374 |
| 63. UI administration.....               | 385 |
| 64. Users and roles directory.....       | 401 |
| 65. Data model administration.....       | 405 |
| 66. Database mapping administration..... | 407 |
| 67. Workflow management.....             | 411 |
| 68. Task scheduler.....                  | 415 |
| 69. Audit trail.....                     | 421 |
| 70. Other.....                           | 423 |

## Distributed Data Delivery (D3)

|                                               |     |
|-----------------------------------------------|-----|
| 71. Introduction to D3.....                   | 426 |
| 72. D3 broadcasts and delivery dataspace..... | 431 |
| 73. D3 JMS Configuration.....                 | 435 |
| 74. D3 administration.....                    | 443 |

# Security Guide

---

75. Security Best Practices.....454

# Developer Guide

---

## Introduction

|                                       |     |
|---------------------------------------|-----|
| 76. Packaging TIBCO EBX® modules..... | 461 |
| 77. Mapping to Java.....              | 467 |
| 78. Tools for Java developers.....    | 473 |
| 79. Terminology changes.....          | 475 |

## Data model

|                                              |     |
|----------------------------------------------|-----|
| 80. Introduction.....                        | 478 |
| 81. Data types.....                          | 481 |
| 82. Tables and relationships.....            | 495 |
| 83. Constraints, triggers and functions..... | 515 |
| 84. Labels and messages.....                 | 533 |
| 85. Additional properties.....               | 539 |
| 86. Data services.....                       | 547 |
| 87. Toolbars.....                            | 549 |
| 88. Workflow model.....                      | 551 |

## User interface

|                                  |     |
|----------------------------------|-----|
| 89. Interface customization..... | 562 |
|----------------------------------|-----|

### *User services*

|                                      |     |
|--------------------------------------|-----|
| 90. Overview.....                    | 565 |
| 91. Quick start.....                 | 569 |
| 92. Implementing a user service..... | 573 |
| 93. Declaring a user service.....    | 587 |
| 94. Development recommendations..... | 593 |

## SOAP data services

|                          |     |
|--------------------------|-----|
| 95. Introduction.....    | 598 |
| 96. WSDL generation..... | 609 |
| 97. SOAP operations..... | 617 |

## REST data services

|                                    |     |
|------------------------------------|-----|
| 98. Introduction.....              | 652 |
| 99. Built-in RESTful services..... | 659 |
| 100. JSON format.....              | 695 |
| 101. REST Toolkit.....             | 721 |



---

# User Guide

---

---

# Introduction

---

## CHAPTER 1

# How TIBCO EBX® works

This chapter contains the following topics:

1. [Product overview](#)
2. [EBX® architecture](#)

## 1.1 Product overview

Master Data Management (MDM) is a way to model, manage and ultimately govern shared data. When data needs to be shared by various IT systems, as well as different business teams, having a single governed version of master data is crucial.

With EBX®, business and IT users can collaborate on a single, unified solution in order to design data models and manage master data content.

EBX® is an MDM software that allows modeling any type of master data and implementing governance using the rich features included, such as collaborative workflows, data authoring, hierarchy management, version control, and role-based security.


An MDM project using EBX® starts with the creation of a *data model*. This is where tables, fields, links and business rules related to the master data are defined. Examples of modeled data include product catalogs, financial hierarchies, lists of suppliers or simple reference tables.

The data model can then be published to make it available to *datasets*, which store the actual master data based on the structure defined in the data model. Datasets are organized and contained within *dataspaces*, containers that isolate updates from one another. Dataspaces allow working on parallel versions of data without the modifications impacting other versions.

*Workflows* are an invaluable feature for performing controlled change management or data approval. They provide the ability to model a step-by-step process involving multiple users, both human and automated.

*Workflow models* detail the tasks to be performed, as well as the parties associated with the tasks. Once a workflow model is published, it can be executed as *data workflows*. Data workflows can notify users of relevant events and outstanding work in a collaborative context.

*Data services* help integrate EBX® with third-party systems (middleware), by allowing external systems to access data in the repository, or to manage dataspace and workflows through web services.



#### See also

[Data modeling](#) [p 22]

[Datasets](#) [p 24]

[Dataspaces](#) [p 26]


[Workflow modeling](#) [p 27]

[Data workflows](#) [p 28]

[Data services](#) [p 29]

## 1.2 EBX® architecture

The following diagram illustrates the EBX® architecture.





## CHAPTER 2

# Using the user interface

This chapter contains the following topics:

1. [Overview](#)
2. [Advanced perspective](#)
3. [Perspectives](#)
4. [User pane](#)
5. [User interface features](#)
6. [Where to find EBX® help](#)

## 2.1 Overview

The general layout of TIBCO EBX® workspaces is entirely customizable by a perspective administrator.

If several customized perspectives have been created, the tiles icon 'Select perspective' allows the user to switch between available perspectives.

The advanced perspective is accessible by default.

See also [UI administration](#) [p 385]

## 2.2 Advanced perspective

By default, the EBX® advanced perspective is available to all users, but its access can be restricted to selected profiles. The view is separated into several general areas, referred to as the following in the documentation:


### Note

The advanced perspective is still accessible to users through explicit selection (for example through a Web component). Unlike other perspectives, it can only be "hidden" in the user interface so that users cannot apply it themselves.

- **Header:** Displays the avatar of the user currently logged in and the perspective selector. Clicking on the user's avatar gives access to the user pane.
- **Menu bar:** The functional categories accessible to the current user.

- **Navigation pane:** Displays context-dependent navigation options. For example: selecting a table in a dataset, or a work item in a workflow.
- **Workspace:** Main context-dependent work area of the interface. For example, the table selected in the navigation pane is displayed in the workspace, or the current work item is executed in the workspace.

The following functional areas are displayed according to the permissions of the current user: *Data*, *Dataspaces*, *Modeling*, *Data Workflow*, *Data Services*, and *Administration*.



## 2.3 Perspectives


The EBX® perspectives are highly configurable views with a target audience. Perspectives offer a simplified user interface to business users and can be assigned to one or more profiles. This view is split into several general areas, referred to as the following in the documentation:

- **Header:** Displays the avatar of the user currently logged in and the perspective selector (when more than one perspective is available). Clicking on the user's avatar gives access to the user pane.
- **Navigation pane:** Displays the hierarchical menu as configured by the perspective administrator. It can be expanded or collapsed to access relevant entities and services related to the user's activity.
- **Workspace:** Main context-dependent work area of the interface.

Perspectives are configured by authorized users. For more information on how to configure a perspective, see [perspective administration](#) [p 386].




Example of a hierarchical menu:



## Favorite perspectives

When more than one perspective is available to a user, it is possible to define one as their favorite perspective so that, when logging in, this perspective will be applied by default. To do so, an icon is available in the perspective selector next to each perspective:

- A full star indicates the favorite perspective. A click on it will remove the favorite perspective.
- An empty star indicates that the associated perspective is not the favorite one. A click on it will set this perspective as the favorite one.



See also [Recommended perspectives](#) [p 397]

## 2.4 User pane

General EBX® features are grouped in the user pane. It can be accessed by clicking on the avatar (or user's initials) in the upper right corner of any page.

The user pane is then displayed with the user avatar and gives access to the profile configuration (according to the user's rights), language selection, density selection and online documentation.

**Attention**

The logout button is located on the user pane.

**Avatar**

An avatar can be defined for each user. The avatar consists in a picture, defined using a URL path; or in two letters (the user's initials by default). The background color is set automatically and cannot be modified. Regarding the image that will be used, it has to be a square format but there is no size limitation.

**Note**

Avatars appear in the user pane, history and workflow interfaces.

The feature is also available through the Java method `UIComponentWriter.addUserAvatarAPI`.

The avatar layout can be customized in the 'Ergonomics and layout' section of the 'Administration' area. It is possible to choose between the display of the avatar only, user name only, or to display both.


**Density**

Users can now choose their display density mode between 'Compact' and 'Comfortable'. The display mode can be modified from the user pane.

## 2.5 User interface features

### ***Resetting the navigation pane width***

After having resized the width of the navigation pane, you can restore it to the default width by hovering over the border and double-clicking.



## 2.6 Where to find EBX® help


In addition to the full standalone product documentation accessible via the [user pane](#) [p 17], help is accessible in various forms within the interface.

### ***Context-sensitive help***


When browsing any workspace in EBX®, context-specific help is available by clicking on the question mark located to the right side of the second header. The corresponding chapter from the product documentation will be displayed.

## Contextual help on elements

When you hover over an element for which contextual help has been defined, a question mark appears. Clicking on the question mark opens a panel with information on the element.



When a permalink to the element is available, a link button appears in the upper right corner of the panel.



## CHAPTER 3

---

# Glossary

This chapter contains the following topics:

1. [Governance](#)
2. [Data modeling](#)
3. [Datasets](#)
4. [Data management life cycle](#)
5. [History](#)
6. [Workflow modeling](#)
7. [Data workflows](#)
8. [Data services](#)
9. [Cross-domain](#)

### 3.1 Governance

#### ***repository***

A back-end storage entity containing all the data managed by TIBCO EBX®. The repository is organized into dataspace.

See also [dataspace](#) [p 26].

#### ***profile***

The generic term for a user or a role. Profiles are used in data workflows and for defining permission rules.

See also [user](#) [p 21], [role](#) [p 22].

Related Java API Profile<sup>API</sup>.

#### ***user***

An entity created in the repository in order for physical users or external systems to authenticate and access EBX®. Users may be assigned roles and have other account information associated with them.

See also [user and roles directory](#) [p 22], [profile](#) [p 21].

Related concept [User and roles directory](#) [p 401].

Related Java API `UserReference`<sup>API</sup>.

## ***role***

A user classification, used for permission rules and data workflows, which can be assigned to users. Each user may belong to multiple roles.

Whenever a role profile is specified in EBX®, the behavior resulting from that designation is applied to all users that are members of that role. For example, in a workflow model, a role may be specified when defining to whom work items are offered. As a result, all users belonging to that role can receive the same work item offer.

See also [user and roles directory](#) [p 22], [profile](#) [p 21].

Related concept [User and roles directory](#) [p 401].

Related Java API `Role`<sup>API</sup>.

## ***administrator***

A predefined role that has access to the technical administration and configuration of EBX®.

## ***user and roles directory***

A directory defining the methods available for authentication when accessing the repository, all available roles, and the users authorized to access the repository with their role assignments.

See also [user](#) [p 21], [role](#) [p 22].

Related concept [User and roles directory](#) [p 401].

Related Java API `Directory`<sup>API</sup>, `DirectoryHandler`<sup>API</sup>.

## ***user session***

A repository access context that is associated with a user after being authenticated against the user and roles directory.

Related concept [User and roles directory](#) [p 401].

Related Java API `Session`<sup>API</sup>.

# 3.2 Data modeling

Main documentation section [Data models](#) [p 32]

## ***data model***

A structural definition of the data to be managed in the EBX® repository. A data model includes detailed descriptions of all included data, in terms of organization, data types, and semantic relationships. The purpose of data models is to define the structure and characteristics of datasets, which are instances of data models that contain the data being managed by the repository.

See also [dataset](#) [p 24].

Related concept [Data models](#) [p 32].

**field**

A data model element that is defined with a name and a simple datatype. A field can be included in the data model directly or as a column of a table. In EBX®, fields can be assigned basic constraints, such as length and size, as well as more complex validation rules involving computations. Automated value assignment using field inheritance or computations based on other data can also be defined for fields. Aggregated lists can be created by setting the cardinality of a field to allow multiple values in the same record. Fields can be arranged into groups to facilitate structural organization in the data model.

By default, fields are denoted by the icon .

See also [record](#) [p 24], [group](#) [p 23], [table \(in data model\)](#) [p 23], [validation rule](#) [p 24], [inheritance](#) [p 25].

Related concepts [Structure elements properties](#) [p 51], [Controls on data fields](#) [p 65].

Related Java API SchemaNode<sup>API</sup>.

The former name (prior to version 5) of "field" was "attribute".

**primary key**

A field or a composition of multiple fields used to uniquely identify the records in a table.

Primary keys are denoted by the icon .

Related concept [Tables definition](#) [p 495].

**foreign key**

A field or a composition of multiple fields in one table whose field values correspond to the primary keys of another table. Foreign keys are used to reference records in one table from another table.

Foreign keys are denoted by the icon .

See also [primary key](#) [p 23].

Related concept [Foreign key](#) [p 500].

**table (in data model)**

A data model element that is comprised of fields and/or groups of fields. Every table must define at least one field to act as the unique identifier, or primary key, of records. A table in a data model can be used to create a reusable type based on the table's structure, which can then be used to create other elements of the same structure in the data model.

Tables are represented by the icon .

See also [record](#) [p 24], [primary key](#) [p 23], [reusable type](#) [p 24].

**group**

A classification entity used to facilitate the organization of a data model. A group can be used to collect fields, other groups, and tables. If a group contains tables, the group cannot be included within another table, as the constraint that tables cannot be nested must be respected. A group can be used to create a reusable type based on the group's structure, which can then be used to create other elements of the same structure in the data model.

Groups are represented by the icon .

See also [reusable type](#) [p 24].

Related Java API SchemaNode<sup>API</sup>.

### ***reusable type***

A shared simple or complex type definition that can be used to define other elements in the data model.

### ***validation rule***

An acceptance criterion defined on a field or a table. Data is considered invalid if it does not comply with all imposed validation rules.

The former name (prior to version 5) of "validation rule" was "constraint".

### ***data model assistant (DMA)***

The EBX® user interface includes a tool that aids the implementation of data models. It allows defining the structure of data models, creating and editing elements, as well as configuring and publishing data models.

See also [Data models](#) [p 32].

## **3.3 Datasets**

Main documentation section [Datasets](#) [p 106]

### ***record***

A set of field values in a table, uniquely identified by a primary key. A record is a row in the table. Each record follows the data structure defined in the data model. The data model drives the data types and cardinality of the fields found in records.

See also [table \(in dataset\)](#) [p 24], [primary key](#) [p 23].

The former name (prior to version 5) of "record" was "occurrence".

### ***table (in dataset)***

A set of records (rows) of the same structure containing data. Each record is uniquely identified by its primary key.

Tables are represented by the icon .

See also [record](#) [p 24], [primary key](#) [p 23].

### ***dataset***

A data-containing instance of a data model. The structure and behavior of a dataset are based upon the definitions provided by the data model that it is implementing. Depending on its data model, a dataset contains data in the form of tables, groups, and fields.

Datasets are represented by the icon .

See also [table \(in dataset\)](#) [p 24], [field](#) [p 23], [group](#) [p 23], [views](#) [p 25].

Related concept [Datasets](#) [p 106].

The former name (prior to version 5) of "dataset" was "adaptation instance".



## ***inheritance***

A mechanism by which data can be acquired by default by one entity from another entity. In EBX®, there are two types of inheritance: dataset inheritance and field inheritance.

When enabled, dataset inheritance allows a child dataset to acquire default data values from its parent dataset. This feature can be useful when designing a data model where data declared in a parent scope will be used with the same value by default in nested child scopes. Values that are inherited from the parent can be overridden by the child. By default, dataset inheritance is disabled. It can be enabled during the data model definition.

Inheritance from the parent dataset is represented by the icon .

Field inheritance is defined in the data model to automatically fetch a field value from a record in another table.

Inherited fields are represented by the icon .

Related concept [Inheritance and value resolution](#) [p 268].

## ***views***

A customizable display configuration that may be applied to viewing tables. A view can be defined for a given user or role, in order to specify whether records are displayed in a tabular or hierarchical format, as well as to set record filtering criteria.

The hierarchical view type offers a tree-based representation of the data in a table. Nodes in the tree can represent either field values or records. A hierarchical view can be useful for showing the relationships between the model data. When creating a view that uses the hierarchical format, dimensions can be selected to determine the structural representation of data. In a hierarchical view, it is possible to navigate through recursive relationships, as well as between multiple tables using foreign key relationships.

### **See also**

[Views](#) [p 113]

[Hierarchies](#) [p 115]

## ***recommended view***

A recommended view can be defined by the dataset owner for each target profile. When a user logs in with no view specified, their recommended view (if any) is applied. Otherwise, the default view is applied.

The 'Manage recommended views' action allows defining assignment rules for recommended views depending on users and roles.

Related concept [Recommended views](#) [p 116].

## ***favorite view***

When displaying a table, the user can choose to define the current as their favorite view through the 'Manage views' sub-menu.

Once it has been set as the favorite, the view will be automatically applied each time this user accesses the table.

Related concept [Manage views](#) [p 117].

## 3.4 Data management life cycle

Main documentation section [Dataspaces](#) [p 88]

### ***dataspace***

A container entity comprised of datasets. It is used to isolate different versions of datasets or to organize them.

Child dataspaces may be created based on a given parent dataspace, initialized with the state of the parent. Datasets can then be modified in the child dataspaces in isolation from their parent dataspace as well as each other. The child dataspaces can later be merged back into their parent dataspace or compared against other dataspaces.

See also [inheritance](#) [p 25], [repository](#) [p 21], [dataspace merge](#) [p 26].

Related concept [Dataspaces](#) [p 88].

The former name (prior to version 5) of "dataspace" was "branch" or "snapshot".

### ***reference dataspace***

The root ancestor dataspace of all dataspaces in the EBX® repository. As every dataspace merge must consist of a child merging into its parent, the reference dataspace is never eligible to be merged into another dataspace.

See also [dataspace](#) [p 26], [dataspace merge](#) [p 26], [repository](#) [p 21].

### ***dataspace merge***

The integration of the changes made in a child dataspace since its creation into its parent dataspace. The child dataspace is closed after the merge has completed successfully. To perform a merge, all the differences identified between the source dataspace and the target dataspace must be reviewed, and conflicts must be resolved. For example, if an element has been modified in both the parent and child dataspace since the creation of the child dataspace, the conflict must be resolved manually by deciding which version of the element should be kept as the result of the merge.

Related concept [Merge](#) [p 96].

### ***snapshot***

A static copy of a dataspace that captures its state and all of its content at a given point in time for reference purposes. A snapshot may be viewed, exported, and compared to other dataspaces, but it can never be modified directly.

Snapshots are represented by the icon .

Related concept [Snapshot](#) [p 101]

The former name (prior to version 5) of "snapshot" was "version" or "home".

## 3.5 History

Main documentation section [History](#) [p 249]

**historization**

A mechanism that can be enabled at the table level to track modifications in the repository. Two history views are available when historization is activated: table history view and transaction history view. In all history views, most standard features for tables, such as export, comparison, and filtering, are available.

Activation of historization requires the configuration of a history profile. The historization of tables is not enabled by default.

See also [table history view](#) [p 27], [transaction history view](#) [p 27], [history profile](#) [p 27].

**history profile**

A set of preferences that specify which dataspace should have their modifications recorded in the table history, and whether transactions should fail if historization is unavailable.

See also [history profile](#) [p 27].

**table history view**

A view containing a trace of all modifications that are made in a given table, including record creations, updates, and deletions. Each entry includes transactional information, such as a timestamp and the user performing the action, as well as the data at the conclusion of the transaction. This information can also be consulted at a record or dataset level.

Related technical reference [History](#) [p 249].

**transaction history view**

A view displaying the technical and authentication data of transactions, either globally at the repository level, or at the dataspace level. As a single transaction can perform multiple actions and affect multiple tables in one or more datasets, this view shows all the modifications that have occurred across the given scope for each transaction.

Related technical reference [History](#) [p 249].

## 3.6 Workflow modeling

Main documentation section [Workflow models](#) [p 130]

**workflow model**

A procedural definition of operations to be performed on data. A workflow model describes the complete path that the data must follow in order to be processed, including its states and associated actions to be taken by human users and automated scripts.

Related concept [Workflow models](#) [p 130].

The former name (prior to version 5) of "workflow model" was "workflow definition".

Workflow models are represented by the icon .

***script task***

A data workflow task performed by an automated process, with no human intervention. Common script tasks include dataspace creation, dataspace merges, and snapshot creation.

Script tasks are represented by the icon .

See also [workflow model](#) [p 27].

***user task***

A data workflow task that is made up of one or more work items performed concurrently by human users. User task work items are offered or assigned to users, depending on the workflow model. The progression of a data workflow beyond a user task depends on the satisfaction of the task termination criteria defined in the workflow model.

User tasks are represented by the icon .

See also [workflow model](#) [p 27].

***workflow condition***

A decision step in a data workflow. A data workflow condition describes the criteria used to decide which step will be executed next.

Workflow conditions are represented by the icon .

***sub-workflow invocation***

A step in a data workflow that pauses the current data workflow and launches one or more other data workflows. If multiple sub-workflows are invoked by the same sub-workflow invocation step, they will be executed concurrently, in parallel.

***wait task***

A step in a data workflow that pauses the current workflow and waits for a specific event. When the event is received, the workflow is resumed and automatically goes to the next step.

***data context***

A set of data that may be shared between steps throughout a data workflow to ensure continuity between steps.

## 3.7 Data workflows

Main documentation section [Data workflows](#) [p 158]

***workflow publication***

An instance of a workflow model that has been made available for execution to users with the appropriate permissions.

The former name (prior to version 5) of "workflow publication" was "workflow".

**data workflow**

An executed instance of a workflow model, which runs the data processing steps that are defined in the model, including user tasks, script tasks, and conditions.

See also [workflow model](#) [p 27].

Related concept [Data workflows](#) [p 158].

The former name (prior to version 5) of "data workflow" was "workflow instance".

**work list**

A list of all published data workflows that the current user has the permissions to view. Users with the permissions to launch data workflows do so from their 'Work List'. All outstanding work items requiring action from the user appear under their published workflows in the work list. Additionally, if the user is the administrator of data workflows, they are able to view the state of execution of those data workflows in their 'Work List', and may intervene if necessary.

**work item**

An action that must be performed by a human user as a part of a user task.

Allocated work items are represented by the icon .

See also [user task](#) [p 28].

**token**

Tokens are used during data workflow management, and are visible to repository administrators.

## 3.8 Data services

Main documentation section [Data services](#) [p 178]

**data service**

EBX® shares master data according to the [Service-oriented architecture](#) (SOA) by using XML web services. Since all data services are generated directly from models or built-in services they can be used to access part of the features available from the user interface.

Data services offer:

- a WSDL model-driven and built-in generator to build a communication interface. It can be produced through the user interface or the HTTP(S) connector for a client application. XML messages are communicated to the EBX® entry point.
- a SOAP connector or entry point component for SOAP messages which allows external systems interacting with the EBX® repository. This connector responds to requests coming from the WSDL produced by EBX®. This component accepts all SOAP XML messages corresponding to the EBX® WSDL generator.
- A RESTful connector, or entry point for the select operations, allows external systems interrogating the EBX® repository. After authenticating, it accepts the request defined in the URL and executes it according to the permissions of the authenticated user.

## ***lineage***

A mechanism by which access rights profiles are implemented for data services. Access rights profiles are then used to access data via WSDL interfaces.

Related concept: [Generating a WSDL for lineage](#) [p 183].

## **3.9 Cross-domain**

### ***node***

A node is an element of a tree view or a graph. In EBX®, 'Node' can carry several meanings depending on the context of use:

- In the [workflow model](#) [p 27] context, a node is a workflow step or condition.
- In the [data model](#) [p 22] context, a node is a group, a table or a field.
- In the [hierarchy](#) [p 25] context, a node represents a value of a dimension.
- In an [adaptation tree](#) [p 25], a node is a dataset.
- In a [dataset](#) [p 24], a node is the node of the data model evaluated in the context of the dataset or the record.

---

# Data models

---

## CHAPTER 4

# Introduction to data models

This chapter contains the following topics:

1. [Overview](#)
2. [Using the Data Models area user interface](#)

## 4.1 Overview

### ***What is a data model?***

The first step towards managing data in TIBCO EBX® is to develop a data model. The purpose of a data model is to provide the detailed structural definition of the data that will be managed in the repository, in terms of organization, data types, and semantic relationships.

In order to implement a data model in the repository, you will first create a new data model, then define the details of the structure of its component table, field, and group elements, as well as their behavior and properties. When you have completed the entry or import of your data model structure in the repository, you will publish it to make it available for use by datasets. Once you have a publication of your data model, you and other users can create datasets based upon it to contain the data that is managed by the EBX® repository.

### ***Basic concepts used in data modeling***

A basic understanding of the following terms is necessary to proceed with the creation of data models:

- [field](#) [p 23]
- [primary key](#) [p 23]
- [foreign key](#) [p 23]
- [table \(in data model\)](#) [p 23]
- [group](#) [p 23]
- [reusable type](#) [p 24]
- [validation rule](#) [p 24]




## 4.2 Using the Data Models area user interface

### *Navigating within the Data Model Assistant*

Data models can be created, edited or imported, and published in the **Data Models** area of the user interface. The EBX® data model assistant (DMA) facilitates the development of data models.

**Note**

This area is available only to authorized users in the 'Advanced perspective'.



The navigation pane is organized into the following sections:

|                             |                                                                                                                                                                                                                                                                                  |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Configuration</b>        | The technical configuration of the data model.                                                                                                                                                                                                                                   |
| <b>Global properties</b>    | Defines the global properties of the data model.                                                                                                                                                                                                                                 |
| <b>Included data models</b> | Defines the data models included in the current model. All types defined in included data models can be reused in the current model.                                                                                                                                             |
| <b>Java bindings</b>        | The bindings specify what Java types have to be generated from the model.                                                                                                                                                                                                        |
| <b>Component library</b>    | Defines the Java components available in the model. These provide programmatic features that will be available for the model, such as programmatic constraints, functions, and UI beans.                                                                                         |
| <b>Toolbars</b>             | The toolbars available to use in the data model.                                                                                                                                                                                                                                 |
| <b>Ajax components</b>      | Defines the available Ajax components in the model.                                                                                                                                                                                                                              |
| <b>User services</b>        | Declares the user services using the API available before release 5.8.0. From release 5.8.0, it is advised to use the new UserService API (these services are directly registered through the Java API, hence no declaration is required for them in the data model assistant).. |
| <b>Add-ons</b>              | Specifies which add-ons are used by the data model. These add-ons will have the capacity to enrich the current data model after the publication by adding properties and constraints to the elements of the data model.                                                          |
| <b>Data services</b>        | Specifies the WSDL operations' suffixes that allow to refer to a table in the data service operations using a unique name instead of its path.                                                                                                                                   |
| <b>Replications</b>         | This table defines the replication units of the data model. A replication unit allows the replication of a source table in the relational database, so that external systems can access this data by means of plain SQL requests and views.                                      |
| <b>Data structure</b>       | The structure of the data model. Defines the relationship between the elements of the data model and provides access to the definition of each element.                                                                                                                          |

|                                    |                                                                    |
|------------------------------------|--------------------------------------------------------------------|
| <b>Simple data types</b>           | Simple reusable types defined in the current data model.           |
| <b>Complex data types</b>          | Complex reusable types defined in the current data model.          |
| <b>Included simple data types</b>  | Simple reusable types defined in an included external data model.  |
| <b>Included complex data types</b> | Complex reusable types defined in an included external data model. |

**See also**

[\*Implementing the data model structure\*](#) [p 45]

[\*Configuring the data model\*](#) [p 39]

[\*Reusable types\*](#) [p 47]

**Data model element icons**

 [field](#) [p 23]

 [primary key](#) [p 23]

 [foreign key](#) [p 23]

 [table](#) [p 23]

 [group](#) [p 23]

**Related concepts**

[\*Dataspaces\*](#) [p 88]

[\*Datasets\*](#) [p 106]



## CHAPTER 5

# Creating a data model

This chapter contains the following topics:

1. [Creating a new data model](#)
2. [Selecting a data model type](#)

## 5.1 Creating a new data model

To create a new data model, click the **Create** button in the pop-up, and follow through the wizard.

## 5.2 Selecting a data model type

If you are a user with the 'Administrator' role, you must choose the type of data model you are creating, *semantic* or *relational*, in the first step of the data model creation wizard.

### ***Semantic models***

Semantic models enable the full use of data management features, such as life cycle management using dataspace, provided by TIBCO EBX®. This is the default type of data model.

### ***Relational models***

Relational models are used when the tables created from a data model will eventually be mapped to a relational database management system (RDBMS). The primary benefit of using a relational model is the ability to query the tables created from the data model using external SQL requests. However, employing a relational model results in the loss of functionalities in EBX® such as inheritance, multi-valued fields, and the advanced life cycle management provided by dataspace.

#### **Note**

A relational data model can only be used by a single dataset, and the dataspace containing the dataset must also be declared as being relational.

#### **See also**

[Relational mode](#) [p 243]

[Dataspace](#) [p 88]



## CHAPTER 6

# Configuring the data model

This chapter contains the following topics:

1. [Information associated with a data model](#)
2. [Permissions](#)
3. [Data model properties](#)
4. [Included data models](#)
5. [Data services](#)
6. [Replication of data to relational tables](#)
7. [Add-ons used by the data model](#)

## 6.1 Information associated with a data model

To view and edit the owner and documentation of your data model, select 'Information' from the [data model 'Actions'](#) [p 33] menu for your data model in the navigation pane.

**Note**

This area is available only to authorized users in the 'Advanced perspective'.

|                                |                                                                                                                           |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Unique name</b>             | The unique name of the data model. This name cannot be modified once the data model has been created.                     |
| <b>Owner</b>                   | Specifies the data model owner, who will have permission to edit the data model's information and define its permissions. |
| <b>Localized documentation</b> | Localized labels and descriptions for the data model.                                                                     |

## 6.2 Permissions

To define the user permissions on your data model, select 'Permissions' from the [data model 'Actions'](#) [p 33] menu for your data model in the navigation pane.

The configuration of the permissions of a data model are identical to the options for the permissions of a dataset, as explained in [Permissions](#) [p 121].

## 6.3 Data model properties

In the navigation pane, under Configuration > Data model properties, you can access the following technical properties:

|                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Module name</b>                   | Defines the module that contains the resources that will be used by this data model. This is also the target module used by the data model publication if publishing to a module.                                                                                                                                                                                                                                                                          |
| <b>Module path</b>                   | Physical location of the module on the server's file system.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Sources locations</b>             | The source paths used when configuring Java components in the 'Component library'. If a path is relative, it will be resolved using the 'Module path' as the base path.                                                                                                                                                                                                                                                                                    |
| <b>Publication mode</b>              | <p>Whether to publish the data model as an XML Schema Document within a module or as a publication completely embedded in the TIBCO EBX® repository. Embedded data models offer additional functionality such as versioning and rollback of publications.</p> <p>See <a href="#">Publication modes</a> [p 83] for more information.</p> <p><b>Model path in module:</b> Defines the target file for the data model generation. It must start with '/'.</p> |
| <b>Dataset inheritance</b>           | <p>Specifies whether dataset inheritance is enabled for this data model. Dataset inheritance is disabled by default.</p> <p>See <a href="#">Dataset inheritance</a> [p 125] for more information.</p>                                                                                                                                                                                                                                                      |
| <b>Documentation</b>                 | <p>Documentation of the data model defined by a Java class. This Java class can programmatically specify labels and descriptions for the elements of the data model. The labels and descriptions defined in this Java class are displayed in associated datasets in preference to the ones defined locally on an element.</p> <p>See <a href="#">Dynamic labels and descriptions</a> [p 534] for more information.</p>                                     |
| <b>Special extensions</b>            | Access permissions defined by programmatic rules in a Java class.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Disable auto-increment checks</b> | <p>Specifies whether to disable if the check of an auto-incremented field value in associated datasets regarding to the "max value" found in the table being updated.</p> <p>See <a href="#">Auto-incremented values</a> [p 531] for more information.</p>                                                                                                                                                                                                 |



**Enable user services (old API)**

Specifies if user services using the API available before release 5.8.0 can be declared. If 'No', the section 'Configuration > User services' is not displayed (except if at least service has been already declared in this section). From release 5.8.0, it is advised to use the new UserService Java API (these services are directly registered through the Java API, hence no declaration is required in the data model assistant).

See `UserServiceDeclarationAPI` for more information.

## 6.4 Included data models

You can use data types in the current model that are defined in another data model by adding an entry for the other data model in the table under Configuration > Included data models.

When you access the record of an included model in this table, you will find technical information about the model under the **Information** tab. As an included data model could eventually have validation errors, for example, due to deleted Java resources, this view will provide information regarding those issues.

It is only possible to include data models that have no validation errors and have been defined and published as an embedded data model or packaged in a module.

The names of data types must be unique across both locally defined and included type definitions. That is, included data types must not have names that coincide with those of data types defined in the current data model or other included data models.

See also [Including external data models](#) [p 493]

## 6.5 Data services

It is possible to refer to tables in Data Service operations using unique names instead of their paths by defining suffixes for WSDL operations. A WSDL suffix is the association between a table path and a name.

To define a WSDL suffix through the user interface, create a new record in the 'Data services' table under the data model configuration in the navigation pane. A record of this table defines the following properties:

**Table path**

Specifies the path of the table in the current data model that is to be referred by the WSDL operation suffix.

**WSDL operation suffix**

This name is used to suffix all the operation names of the concerned table. If undefined for a given table, the last element of the table path is used instead. This name must be unique in the context of this data model.

See also [Data services](#) [p 547]

## 6.6 Replication of data to relational tables

In any data model, it is possible to define *replication units* for data in the repository to be mirrored to dedicated relational tables. These relational tables then enable direct access to the data by SQL requests and views.

To define a replication unit through the user interface, create a new record in the 'Replications' table under the data model configuration in the navigation pane. Each replication unit record is associated

with a particular dataset in a given dataspace. A single replication unit can cover multiple tables, as long as they are in the same dataset. A replication unit defines the following information:

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>             | Name of the replication unit. This name identifies a replication unit in the current data model. It must be unique.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Dataspace</b>        | Specifies the dataspace relevant to this replication unit. It cannot be a snapshot or a relational dataspace.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Dataset</b>          | Specifies the dataset relevant to this replication unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Refresh policy</b>   | <p>Specifies the data synchronization policy. The possible policies are:</p> <ul style="list-style-type: none"> <li>• <b>On commit:</b> The replicated table content in the database is always up to date with respect to its source table. Every transaction that updates the EBX® source table triggers the corresponding insert, update, and delete statements on the replicated table.</li> <li>• <b>On demand:</b> The replicated table in the database is only updated when an explicit refresh operation is performed.</li> </ul> |
| <b>Tables</b>           | <p>Specifies the tables in the data model to be replicated in the database.</p> <p><b>Table path:</b> Specifies the path of the table in the current data model that is to be replicated to the database.</p> <p><b>Table name in database:</b> Specifies the name of the table in the database to which the data will be replicated. This name must be unique amongst all replications units.</p>                                                                                                                                       |
| <b>Aggregated lists</b> | <p>Specifies the properties of the aggregated lists in the table that are replicated in the database.</p> <p><b>Path:</b> Specifies the path of the aggregated list in the table that is to be replicated to the database.</p> <p><b>Table name in database:</b> Specifies the name of the table in the database to which the data of the aggregated list will be replicated. This name must be unique amongst all replications units.</p>                                                                                               |

See also [Replication](#) [p 257]

## 6.7 Add-ons used by the data model

On any data model, it is possible to specify the *add-ons* used by the current data model. These add-ons will have the capacity to enrich the current data model after the publication by adding properties and constraints to the data model elements.

To define an add-on to be used by the data model through the user interface, create a new record in the 'Add-ons' table under the data model configuration in the navigation pane. A record of this table defines the following properties:

|                  |                                                                                         |
|------------------|-----------------------------------------------------------------------------------------|
| <b>Name</b>      | Add-on public name.                                                                     |
| <b>Version</b>   | Add-on version.                                                                         |
| <b>Activated</b> | Indicates if the add-on is activated. The add-on must be activated in order to be used. |

## CHAPTER 7

# Implementing the data model structure

To work with the structural definition of your data model, select the data model you are working with in the navigation pane.

You can then access the structure of your data model in the navigation pane under 'Data structure', to define the structure of fields, groups, and tables.

This chapter contains the following topics:

1. [Common actions and properties](#)
2. [Reusable types](#)
3. [Data model element creation details](#)
4. [Modifying existing elements](#)

## 7.1 Common actions and properties

### *Adding elements to the data model*

The following elements are available to describe the structure of your data model:

- fields
- groups
- tables
- primary keys
- foreign keys
- associations

Add a new element relative to any existing element in the data structure by clicking the down arrow ▼ to the right of the existing entry, and selecting an element creation option from the menu. Depending on whether the existing element is a field, group, or table, you have the choice of creating the new

element as a child of the existing element, or before or after the existing element at the same level. You can then follow the element creation wizard to create the new element.

**Note**

The element `root` is always added upon data model creation. If this element must be renamed, it can be deleted and recreated with a new name.

## Names, labels, descriptions, and information

Whenever you create a new element in your data model, you must provide a name for the field 'Name' that is unique in its level of the data structure. This name is assigned once the element is created and cannot be changed subsequently.

You have the option to provide localized user-friendly labels to be displayed in the user interface instead of the unique name of the element, as well as brief localized descriptions of the element. Unlike the unique name, the labels and descriptions are modifiable after creation. According to the language preference of each user, TIBCO EBX® will display the corresponding localized label and description of the element.

## Deleting elements of the data model

Any element can be deleted from the data structure using the down arrow ▼ corresponding to its entry.

When deleting a group or table that is not using a reusable type, the deletion is performed recursively, removing all its nested elements.

## Duplicating existing elements

To duplicate an element, click the down arrow ▼ corresponding to its entry. You must provide a name for the duplicated element that is unique at its level of the data structure. All other properties are copied from the source element.

The duplicated element is added to the data model at the same level as the element from which it was copied, appended after the existing elements. If you are duplicating a table or group containing other elements, all nested elements are copied with their original names.

**Note**

If you duplicate a primary key field, the properties of the field are maintained, but the new field is not automatically added to the primary key.

## Moving elements

To reorder an element within its current level of the data structure, click the down arrow ▼ corresponding to its entry and select 'Move'. Then, select the left-arrow button corresponding to the field *before which* you want to move the current element.

**Note**

It is not possible to move an element to a position outside of its level in the data structure.

## 7.2 Reusable types

Reusable types are shared element definitions that are created once and can be reused in different places in the data model.

### Note

If you modify the definition of a reusable type in the 'Simple data types' or 'Complex data types' section, you will modify the structure of all elements based on that reusable type. The structure of a groups or table using a reusable type is shown as read-only. To edit the structure of the associated reusable type, you have to access the type from the 'Simple data types' or 'Complex data types' section.

### ***Defining a reusable type***

From the down arrow ▼ menu of 'Simple data types' and 'Complex data types' in the navigation pane, you can define simple and complex reusable types that will be available for creating more elements which share the same structural definition and properties. Alternatively, you can convert existing tables and groups into reusable types using their corresponding down arrow ▼ menus.

It is possible to see the elements that are using a reusable type by selecting 'References to this type' on the specific page of each data type, under 'Simple data types' and 'Complex data types' in the navigation pane. A table then displays all elements that are based on this type. If a data type is not used by any elements, you can select the 'Delete type' from its down arrow ▼ menu to delete the reusable type.

### ***Using a reusable type***

The structure of new elements can be defined using reusable types. To do so, select an existing reusable type in the element creation form. The created element will then share the type definition of the reusable type.

### ***Including data types defined in other data models***

You can also share reusable types between multiple data models. By configuring the inclusion of an external data model, you can use the data types defined in that data model to create new elements in the data structure the same way as using locally defined reusable types.

### Note

As the names of data types must be unique across all locally defined as well as all included types, you cannot create new reusable types with the same name as a data type in an included data model. Similarly, you cannot include an external data model that defines a data type with the same name as a locally defined reusable type or a data type in another included data model.

Included data types appear in the sections 'Included simple data types' and 'Included complex data types' in the navigation panel. You can view the details of these included reusable types; however, they can only be edited locally in their original data models.

See [Included data models](#) [p 41] for more information.

## 7.3 Data model element creation details

### ***Creating fields***

When creating a new field, you must select its data type, which will define the data type of the values based upon this field. The data type of the field cannot be changed once the field has been created.

While creating a field, it is also possible to designate it as a foreign key, a mandatory field, and, if created under a table, a primary key.

### ***Creating tables***

While creating a table, you have the option to create the new table based on an existing reusable type. See [Reusable types](#) [p 47] for more information.

Every table requires specifying at least one primary key field, which you can create as a child element of the table from the navigation pane.

### ***Creating groups***

While creating a group, you have the option to create the new group based on an existing reusable type. See [Reusable types](#) [p 47] for more information.

### ***Creating primary key fields***

At least one primary key is required for every table. You can create a primary key field for a table by creating it as a child element under the table's entry in the 'Data structure' tree.

Besides creating a new field directly as a primary key, you can add any existing child field of a table to the definition of its primary key on the 'Primary key' tab of the table's 'Advanced properties'.

### ***Creating or defining foreign key fields***

Foreign key fields have the data type 'String'. You can create a foreign key field for a table by creating it as a child element under the table's entry in the 'Data structure' tree. You can also convert an existing field of type 'String' into a foreign key. To convert an existing field of type 'String' into a foreign key, enable 'Foreign key constraint' in the field's 'Advanced controls' and define the associated parameters.

Whether creating a foreign key directly or from an existing field, you must define the table that contains the records to be referenced.

### ***Creating associations***

An association allows defining semantic links between tables. You can create an association by creating it as a child element under the table's entry in the 'Data structure' tree and by selecting 'association' in the form for creating a new element. An association can only be defined inside a table. It is not possible to convert an existing field to an association.

When creating an association, you must specify the type of association. Several options are available:

- Inverse relationship of a *foreign key*. In this case, the association element is defined in a *source table* and refers to a *target table*. It is the counterpart of the foreign key field, which is defined in the target table and refers back the source table. You must define the foreign key that references the parent table of the association.



- Over a *link table*. In this case, the association element is defined in a *source table* and refers to a *target table* that is inferred from a *link table*. This link table defines two foreign keys: one referring to the source table and another one referring to the target table. The primary key of the link table must also refer to auto-incremented fields and/or the foreign key to the source or target table of the association. You must define the link table and these two foreign keys.
- Using an *XPath predicate*. In this case, the association element is defined in a *source table* and refers to a *target table* that is specified using a *path*. An *XPath expression* is also defined to specify the criteria used to associate a record of the current table to records of the target table. You must define the target table and an XPath expression.

In all types of association, we call *associated records* the records in the target table that are semantically linked to records in the source table.

Once you have created an association, you can specify additional properties. For an association, it is then possible to:

- Filter associated records by specifying an additional XPath filter. It is only possible to use fields from the source and the target table when defining an XPath filter. That is, if it is an association other a link table it is not possible to use fields of the link table in the XPath filter. You can use the available wizard to select the fields that you want to use in your XPath filter.
- Configure a tabular view to define the fields that must be displayed in the associated table. It is not possible to configure or modify an existing tabular view if the target table of the association does not exist. If a tabular view is not defined, all columns that a user is allowed to view according to the granted access rights are displayed.
- Define how associated records are to be rendered in forms. You can specify that associated records are to be rendered either directly in the form or in a specific tab. By default, associated records are rendered in the form at the same position of the association in the parent table.
- Hide/show associated records in data service 'select' operation. By default associated records are hidden in data service 'select' operation.
- Specify the minimum and maximum numbers of associated records that are required. In associated datasets, a validation message of the specified severity is added if an association does not comply with the required minimum or the maximum numbers of associated records. By default, the required minimum and the maximum numbers of associated records are not restricted.
- Add validation constraints using XPath predicates to restrict associated records. It is only possible to use fields from the source and the target table when defining an XPath predicate. That is, if it is an association over a link table it is not possible to use fields of the link table in the XPath predicate. You can use the available wizard to select the fields that you want to use in your XPath predicate. In associated datasets, a validation message of the specified severity is added when an associated record does not comply with the specified constraint.

## 7.4 Modifying existing elements

### ***Removing a field from the primary key***

Any field that belongs to the primary key can be removed from the primary key on the 'Primary key' tab of the table's 'Advanced properties'.

See [primary key](#) [p 23] in the glossary.



## CHAPTER 8

---

# Properties of data model elements

After the initial creation of an element, you can set additional properties in order to complete its definition.

See also [Data validation controls on elements](#) [p 65]

This chapter contains the following topics:

1. [Basic element properties](#)
2. [Advanced element properties](#)

## 8.1 Basic element properties

### *Common basic properties*

The following basic properties are shared by several types of elements:

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Information</b>              | Additional non-internationalized information associated with the element.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Minimum number of values</b> | <p>Minimum number of values for an element.</p> <p>As primary keys cannot be multi-valued, they must have this property set to '1' or 'undefined'. The minimum number of values is automatically set to '0' when defining the field as a selection node.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Maximum number of values</b> | <p>Maximum number of values for an element. When set to a value greater than '1', the element becomes multi-valued.</p> <p>As primary keys cannot be multi-valued, they must have this property set to '1' or 'undefined'.</p> <p>For tables, the maximum number of values is automatically set to 'unbounded' upon creation. The maximum number of values is automatically set to '0' when defining the field as a selection node.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Validation rules</b>         | <p>This property is available for tables and fields in tables except Password fields, reusable types, fields in complex reusable types, and selection nodes. Used to define powerful and complex validation rules with the help of the provided XPath 1.0 criteria editor.</p> <p>See <a href="#">Criteria editor</a> [p 291] for more information.</p> <p>This can be useful if the validation of the value depends on complex criteria or on the value of other fields.</p> <p>It is also possible to indicate that a rule defines a verification for a value that is mandatory under certain circumstances. In this case a value is mandatory if the rule is not satisfied. See <a href="#">Constraint on 'null' values</a> [p 523] for more information.</p> <p>Using the associated wizard, you can define localized labels for the validation rule, as well as define a localized message with severity to be displayed if the criteria is not met.</p> <p>When defining the severity of the validation message it is possible to indicate whether an input that would violate a validation rule will be rejected or not when submitting a form. The error management policy is only available on validation rules defined on a field and when the severity is set to 'error'. If the validation rule must remain valid, then</p> |

any input that would violate the rule will be rejected and the values will remain unchanged. If errors are allowed, then any input that would violate the rule will be accepted and the values will change. If not specified, the validation rule always blocks errors upon the form submission by default.

If a validation rule is defined on a table, it will be considered as a 'constraint on table' and each record of the table will be evaluated against it at runtime. See [Constraints on table](#) [p 523] for more information.

---

## ***Basic properties for fields***

The following basic properties are specific to fields:

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default value</b>            | <p>Default value assigned to this field. In new data creation forms, the default value appears automatically in the user input field. The default value must comply with the defined type of the field.</p> <p>See <a href="#">Default value</a> [p 539] for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Conversion error message</b> | <p>Internationalized messages to display to users when they enter a value that is invalid for the data type of this field.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Computation rule</b>         | <p>This property is available for fields in tables, except in reusable types. Defines a rule for computing the value of the field using the provided XPath 1.0 editor.</p> <p>See <a href="#">criteria editor</a> [p 291]</p> <p>This can be useful if the value depends on other values in the same record, but does not require a programmatic computation.</p> <p>The following limitations exist for computation rules:</p> <ul style="list-style-type: none"> <li>• Computation rules can only be defined on simple fields inside a table.</li> <li>• Computation rules cannot be defined on fields of type OResource or Password.</li> <li>• Computation rules cannot be defined on selection nodes and primary key fields.</li> <li>• Computation rules cannot be defined when accessing an element from the validation report.</li> </ul> |

---

## 8.2 Advanced element properties

### *Common advanced properties*

The following advanced properties are shared by several types of elements:

---

**Default view and tools >  
Visibility**

Specifies whether or not this element is shown in the default view of a dataset, in the text search of a dataset or in the data service "select" operation.

- **Model-driven view**

Specifies whether or not the current element is shown in the default tabular view of a table, the default record form of a table, and in the default view of a dataset if the current element is a table. Default dataset view, tabular view and default record form generated from the structure of the data model. If the current element is inside a table, then setting the property to 'Hidden' will hide the element from the default tabular view and default record form of the table without having to define specific access permissions. Current element will still be displayed in the view configuration wizard to be able to create a custom view that displays this element. If the current element is a table, then setting the property to 'Hidden' will hide the table from the default view of a dataset without having to define specific access permissions. This property is ignored if it is set on an element that is neither a table nor in a table.

- **All views**

Specifies whether or not the current element is shown in all views of a table in a dataset. Setting the property to 'Hidden in all views' will hide the element in all views of the table, whether tabular (default tabular view included) or hierarchical, without having to define specific access permissions. The current element will also be hidden in the view configuration wizard. That is, it won't be possible to create a custom view that will display this element. This property is ignored if it is set on an element that is not in a table. This property is not applied on forms. That is, setting the property to 'Hidden in all views' will not hide the element in a record form but only in views.

- **Search tools**

Specifies whether or not the current element is shown in a dataset search tool. Setting the property to 'Hidden in all searches' will hide the element both in the text and typed search tools of a dataset. Setting the property to

'Hidden only in text search' will only hide the element in the text search tool. If this property is not set, the element will be displayed in the text search tool by default. This property is ignored if it is set on an element that is not in a table.

- Data services

Specifies whether or not the current element is shown in the data service select operation. Setting the property to 'Excluded from Data Services' will hide the element in the data service "select" operation. This property is ignored if it is set on an element that is not in a table.

See [Default view](#) [p 541] in the Developer Guide.

---

### Default view and tools > Widget

Defines the widget to be used. A widget is an input component that is displayed in forms in associated datasets. If undefined, a default widget is displayed in associated datasets according to the type and properties of the current element. It is possible to use a built-in widget or a custom widget. A custom widget is defined using a Java API to allow the development of rich user interface components for fields or groups. Built-in and custom widgets cannot be defined on a table or an association. It is forbidden to define both a custom widget and a UI bean. It is forbidden to define on a foreign key field both a custom widget and a combo-box selector.

See `UIWidgetFactory`<sup>API</sup> for more information.

---

### Default view and tools > Combo-box selector

Specifies the name of the published view that will be used in the combo-box selection of the foreign key. A selection button will be displayed at the bottom right corner of the drop-down list. When defining a foreign key, this feature allows accessing an advanced selection view through the 'Selector' button that opens the advanced selection view, from where sorting and searching options can be used. If no published view is defined, the advanced selection view will be disabled. If the path of the referenced table is absolute then only the published views corresponding to this table will be displayed. If the path of the referenced table is relative then all the published views associated with the data model containing the target table will be displayed. This property can only be set if no custom widget is defined.

See [Defining a view for the combo box selector of a foreign key](#) [p 543] in the Developer Guide.

---

### UI bean

#### Attention

From version TIBCO EBX® 5.8.0, it is recommended to use widgets instead of UI Beans. Widgets provide more features than UI Beans, and no further evolution will be made on UI beans. See [\\_widget](#) [p 55] for more information.

This property is available for all elements except tables and associations. Specifies a Java class to customize the user interface associated with this element in a dataset. A UI bean can display the element differently and/or modify its value by extending the `UIBeanEditorAPI` class in the Java API.

---

### Transformation on export

This property is available for fields and for groups that are terminal nodes. Specifies a Java class that defines transformation operations to be performed when exporting an associated dataset as an archive. The input of the transformation is the value of this element.

See `NodeDataTransformerAPI` for more information.

---

### Access properties

Defines the access mode for the current element, that is, if its data can be read and/or written.

- 'Read & Write' corresponds to the mode `RW` in the data model XSD.
- 'Read only' corresponds to the mode `R-` in the data model XSD.
- 'Not a dataset node' corresponds to the mode `CC` in the data model XSD.
- 'Non-terminal node' corresponds to the mode `--` in the data model XSD.

See [Access properties](#) [p 539] in the Developer Guide.

---

### Comparison mode

Defines the comparison mode associated with the element, which controls how its differences are detected in a dataset.

- 'Default' means the element is visible when comparing associated data.
- 'Ignored' implies that no changes will be detected when comparing two versions of modified content (records or datasets).

During a merge, the data values of ignored elements are not merged even if the content has been modified. However, values of ignored data sets or records being created during the operation are merged.

During an archive import, values of ignored elements are not imported when the content has been modified.



However, values of ignored datasets or records being created during the operation are imported.

See [Comparison mode](#) [p 544] in the Developer Guide.

---

#### **Apply last modifications policy**

Defines if this element must be excluded from the service allowing to apply the last modifications that have been performed on a record to the other records of the same table.

- 'Default' means that the last modification on this element can be applied to other records.
- 'Ignored' implies that the last modification on this element cannot be applied to other records. This element will not be visible in the apply last modifications service.

See [Apply last modifications policy](#) [p 544] in the Developer Guide.

---

#### **Node category**

Defines a category for this element. Categories allow controlling the visibility of data in a dataset to users. A node with the category 'Hidden' is hidden by default. Restriction: category specifications other than 'Hidden' do not apply to table record nodes.

See [Categories](#) [p 545] in the Developer Guide.

---

### ***Advanced properties for fields***

The following advanced properties are specific to fields.

#### **Check null input**

Implements the property `osd:checkNullInput`. This property is used to activate and check a constraint on `null` at user input time.

By default, in order to allow for temporarily incomplete input, the check for mandatory elements is not performed upon user input, but only upon dataset validation. If a mandatory element must be checked immediately upon user input, set this property to 'true'.

##### **Note**

A value is considered mandatory if the 'Minimum number of values' property is set to '1' or greater. For terminal elements, mandatory values are only checked in activated datasets. For non-terminal elements, the values are checked regardless of whether the dataset is activated.

See [Constraints, triggers and functions](#) [p 527] in the Developer Guide.

#### **Trim whitespaces**

Trim white spaces

Implements the property `osd:trim`. This property is used to indicate whether leading and trailing white spaces must be trimmed upon user input. If this property is not set, leading and trailing white spaces are removed upon user input.

See [Whitespace handling upon user input](#) [p 528] in the Developer Guide.

## UI bean

See [Common advanced properties](#) [p 55].

## Function (computed value)

This property is available for non-primary key fields. Specifies a Java class that computes the value of this field programmatically. This can be useful if the value of the field depends on other values in the repository, or if the computation of the value needs to retrieve data from a third-party system.

A function can be created by implementing the `ValueFunctionAPI` interface.

## Disable validation

Specifies if the constraints defined on the field must be disabled. This property can only be defined on computed fields. If true, cardinalities, simple and advanced constraints defined on the field won't be checked when validating associated datasets.

## Transformation on export

See [Common advanced properties](#) [p 56].

## Access properties

See [Common advanced properties](#) [p 56].

## Auto-increment

This property is only available for fields of type 'Integer' that are contained in a table. When set, the value of the field is automatically calculated when a new record is created. This can be useful for primary keys, as it generates a unique identifier for each record. Two attributes can be specified:

|                                      |                                                                                                                                                         |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Start value</b>                   | Value with which to begin the auto-increment. If this attribute is not specified, the default value is '1'.                                             |
| <b>Increment step</b>                | Amount the value is incremented based on the previous value of the auto-increment. If this attribute is not specified, the default is value is '1'.     |
| <b>Disable auto-increment checks</b> | Specifies whether to disable the check of the auto-incremented field value in associated datasets against the maximum value in the table being updated. |

Auto-incremented values have the following behavior:

- The computation and allocation of the field value are performed whenever a new record is inserted and the field value is yet undefined.

- No allocation is performed if a programmatic insertion already specifies a non-null value. Consequently, the allocation is not performed for a record insertion in occulting or overwriting modes.
- If an archive import specifies the value, the imported value takes precedence.
- Whenever possible, the newly allocated value is unique in the scope of the repository.

That is, the uniqueness of the allocation spans over all datasets based upon this data model, in any dataspace in the repository. The uniqueness across different dataspace facilitates the merging of child dataspace parent dataspace while reasonably avoiding conflicts when a record's primary key includes the auto-incremented value.

Despite this policy, a specific limitation exists when a mass update transaction assigning specific values is performed concurrently with a transaction that allocates an auto-incremented value on the same field. It is possible that the latter transaction will allocate a value that has already been set in the former transaction, as there is no locking between different dataspace.

See [Auto-incremented values](#) [p 531] in the Developer Guide.

## Default view

See [Common advanced properties](#) [p 54].

## Node category

See [Common advanced properties](#) [p 57].

## Inherited field

Defines a relationship from the current field to a field in another table in order to automatically fetch its field value.

|                       |                                                                                                                                                                                                                                                                                      |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Source record</b>  | A foreign key or white space-separated sequence of foreign keys that leads from the current element to the record from which to inherit this field's value. If this property is not specified, the current record is used as the source record for the inheritance.                  |
| <b>Source element</b> | XPath of the element in the source record from which to inherit this field's value. The source element must be terminal, belong to the record described by 'Source record', and its type must match the type of this field. This property is mandatory when using field inheritance. |

See [inheritance](#) [p 25] in the glossary.

For more information, see also [Inherited fields](#) [p 270].

## Advanced properties for tables

The following advanced properties are specific to tables.

## Table

|                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Primary key</b>                                                    | <p>A list of fields in the table that compose the table's primary key. You can add or remove primary key fields here, as in the 'Data structure' view.</p> <p>Each primary key field is denoted by its absolute XPath notation that starts under the table's root element.</p> <p>If there are several elements in the primary key, the list is white-space delimited. For example, <code>"/name /startDate"</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Presentation</b>                                                   | <p>Specifies how records are displayed in the user interface of this table in a dataset.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b><i>Presentation</i> &gt; Record labeling</b>                       | <p>Defines the fields to provide the default and localized labels for records in the table.</p> <p>Can also specify a Java class to set the label programmatically, or set the label in a hierarchy. This Java class must implement either the <code>UILabelRenderer<sup>API</sup></code> interface or the <code>UILabelRendererForHierarchy<sup>API</sup></code> interface.</p> <p><b>Attention:</b> Access rights defined on associated datasets are not applied when displaying record labels. Fields that are usually hidden due to access rights restrictions will be displayed in labels.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b><i>Presentation</i> &gt; Default rendering for groups in forms</b> | <p>Specifies the default display rendering mode of the groups contained in this table. If nothing is defined here, the default policy set in the Administration area will be used to display groups.</p> <p>See <a href="#">Record form: rendering mode for nodes</a> [p 392] in the Administration Guide.</p> <p><b>Enabled rendering for groups</b></p> <p>Specifies a display rendering mode to be enabled for groups in the table in addition to the modes 'expanded' and 'collapsed', which are always available. Tabs must be enabled on the table to have the option to display groups as tabs. Similarly, links must be enabled to have the option to display groups as links.</p> <p><b>Default rendering for groups</b></p> <p>Specifies the default display rendering mode to use for the groups contained in this table. If a group does not specify a default mode then the default mode defined for this table will be used. Links must be enabled to define the default rendering mode as 'Link'. Select a rendering mode according to network and browser performance. Link mode is lighter</p> |

as its content is not displayed on the same page, whereas the expanded and collapsed modes are heavier.

**Note:** When the tabs view is enabled on a table, any groups that would otherwise have used links are automatically converted to collapsed mode. This is done to avoid the inherent display layout complexities that would arise from links and tabs coexisting in the same user interface.

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Presentation &gt; Specific rendering of forms</b> | <p>Defines a specific rendering for customizing the record form in a dataset.</p> <p>See <code>UIForm<sup>API</sup></code> and <code>UserServiceRecordFormFactory<sup>API</sup></code> for more information.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Toolbars</b>                                      | <p>Defines the toolbars to use in this table.</p> <p>Toolbars can be edited in the <i>Configuration &gt; Toolbars</i> section.</p> <p><b>Tabular view top:</b> Defines the toolbar to use on top of the default table view.</p> <p><b>Tabular view row:</b> Defines the toolbar to use on each row of the default table view.</p> <p><b>Record top:</b> Defines the toolbar to use in the record form.</p> <p><b>Hierarchy top:</b> Defines the toolbar to use in the default hierarchy view of the table.</p> <p>See <a href="#">Toolbars</a> [p 73] for more information.</p> |
| <b>History</b>                                       | <p>Specifies when historization is to be performed, and the level of guarantee requested. The available history profiles can be edited in <b>Administration &gt; History and logs</b>.</p> <p>See <a href="#">History configuration in the repository</a> [p 249] for more information.</p>                                                                                                                                                                                                                                                                                     |
| <b>Indexes</b>                                       | <p>Defines the fields to index in the table. Indexing speeds up table access for requests on the indexed fields. No two indexes can contain the exact same fields.</p> <p><b>Index name:</b> Unique name for this index.</p> <p><b>Fields to index:</b> The fields to index, each represented by an absolute XPath notation starting under the table root element.</p>                                                                                                                                                                                                          |
| <b>Specific filters</b>                              | <p>Defines record display filters on the table.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Actions</b>                                       | <p>Specifies the actions that are allowed on the table in associated datasets. By default, all actions are allowed unless specific access rights are defined in a dataset.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |

## Uniqueness constraints

Indicates which fields or set of fields must be unique across the table.

## Triggers

Specifies Java classes that defines methods to be automatically executed when modifications are performed on the table, such as record creation, updates, deletion, etc.

A built-in trigger for starting data workflows is included by default.

See [Triggers](#) [p 530] in the Developer Guide.

## Access properties

See [Common advanced properties](#) [p 56].

## Default view

See [Common advanced properties](#) [p 54].

## Node category

See [Common advanced properties](#) [p 57].

## *Advanced properties for groups*

The following advanced properties are specific to groups.

### Value container class (JavaBean)

Specifies a Java class to hold the values of the children of this group. The Java class must conform to the JavaBean standard protocol. That is, each child of the group must correspond to a JavaBean property in the class, and all properties must have getter and setter accessors defined.

### UI bean

See [Common advanced properties](#) [p 55].

### Transformation on export

See [Common advanced properties](#) [p 56].

### Access properties

See [Common advanced properties](#) [p 56].

## Default view

---

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Visibility</b>         | See <a href="#">Common advanced properties</a> [p 54].                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <hr/>                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Rendering in forms</b> | <p>Defines the rendering mode of this group. If this property is not set, then the default view for groups specified by the container table will be used. 'Tab' and 'Link' are each only available when the container table enables it.</p> <p><b><i>Tab position</i></b></p> <p>This attribute specifies the position of the tab with respect to all the tabs defined in the model. This position is used for determining tab order. If a position is not specified, the tab will be displayed according to the position of the group in the data model.</p> |

---

## Node category

See [Common advanced properties](#) [p 57].

**Related concepts** [Data validation controls on elements](#) [p 65]





## CHAPTER 9

---

# Data validation controls on elements

After the initial creation of an element, you can set additional controls in order to complete its definition.

See also [Properties of data model elements](#) [p 51]

This chapter contains the following topics:

1. [Simple content validation](#)
2. [Advanced content validation](#)

## 9.1 Simple content validation

Using the following static controls on a field, you can establish simple validation rules on its content. The controls available for a given field are dependent on its data type.

|                                 |                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Fixed length</b>             | The exact number of characters required for this field.                                                                                                                                                                                     |
| <b>Minimum length</b>           | The minimum number of characters allowed for this field.                                                                                                                                                                                    |
| <b>Maximum length</b>           | The maximum number of characters allowed for this field.                                                                                                                                                                                    |
| <b>Pattern</b>                  | A regular expression pattern that the value of the field must match. It is not possible to simultaneously define a pattern for both a field and its data type.                                                                              |
| <b>Decimal places</b>           | The maximum number of decimal places allowed for this field.                                                                                                                                                                                |
| <b>Maximum number of digits</b> | The maximum total number of digits allowed for this integer or decimal field.                                                                                                                                                               |
| <b>Enumeration</b>              | Defines a list of predefined possible values for this field. If enumerations are defined in both a field and its type, then the enumeration of this field in associated datasets is replaced by the intersection of these two enumerations. |
| <b>Greater than [constant]</b>  | Defines the minimum value allowed for this field.                                                                                                                                                                                           |
| <b>Less than [constant]</b>     | Defines the maximum value allowed for this field.                                                                                                                                                                                           |

See [XML schema supported facets](#) [p 515].

## 9.2 Advanced content validation

Using the following dynamic and contextual controls on an element, you can establish advanced validation rules of its content. The controls available for a given element are dependent on the type of element and its data type, if it has one.

See also [Dynamic constraints](#) [p 518]

---

## Foreign key constraint

---

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Table</b>                  | Defines the table referenced by the foreign key. A foreign key references a table in the same dataset by default. It can also reference a table in another dataset in the same dataspace, or a dataset in a different dataspace.                                                                                                                                                                                                                                                                                              |
| <b>Mode</b>                   | Location of the table referenced by the foreign key.<br>'Default': current data model.<br>'Other dataset': different dataset, in the same dataspace.<br>'Other dataspace': dataset in a different dataspace.                                                                                                                                                                                                                                                                                                                  |
| <b>Referenced table</b>       | XPath expression describing the location of the table. For example, <code>/root/MyTable</code> .                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Referenced dataset</b>     | Required if the table is located in another dataset. The unique name of the dataset containing the referenced table.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Referenced dataspace</b>   | Required if the table is located in another dataspace. The unique name of the dataspace containing the referenced table.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Label</b>                  | Defines fields to provide the default and localized labels for records in the table.<br><br>Can also specify a Java class to set the label programmatically if 'XPath expression' is set to 'No'. This Java class must implement the <code>TableRefDisplay<sup>API</sup></code> interface of the Java API.<br><br><b>Attention:</b> Access rights defined on associated datasets are not applied when displaying record labels. Fields that are usually hidden due to access rights restrictions will be displayed in labels. |
| <b>Filter</b>                 | Defines a foreign key filter using an XPath expression.<br><br>Can also specify a Java class that implements the <code>TableRefFilter<sup>API</sup></code> interface of the Java API.                                                                                                                                                                                                                                                                                                                                         |
| <b>Greater than [dynamic]</b> | Defines a field to provide the minimum value allowed for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Less than [dynamic]</b>    | Defines a field to provide the maximum value allowed for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

|                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Fixed length [dynamic]</b>             | Defines a field to provide the exact number of characters required for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Minimum length [dynamic]</b>           | Defines a field to provide the minimum number of characters allowed for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Maximum length [dynamic]</b>           | Defines a field to provide the maximum number of characters allowed for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Excluded values</b>                    | Defines a list of values that are not allowed for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Excluded segment</b>                   | <p>Defines an inclusive range of values that are not allowed for this field.</p> <p><b>Minimum excluded value:</b> Lowest value not allowed for this field.</p> <p><b>Maximum excluded value:</b> Highest value not allowed for this field.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Specific constraint (component)</b>    | Specifies one or more Java classes that implement the <code>Constraint<sup>API</sup></code> interface of the Java API. See <a href="#">Programmatic constraints</a> [p 522] for more information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Specific enumeration (component)</b>   | Specifies a Java class to define an enumeration. The class must define an ordered list of values by implementing the <code>ConstraintEnumeration<sup>API</sup></code> interface of the Java API.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Enumeration filled by another node</b> | Defines the possible values of this enumeration using a reference to another list or enumeration element.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Dataspace set configuration</b>        | <p>Define the dataspace identifiers that can be referenced by a field of the type Dataspace identifier (osd:dataspaceKey). If a configuration is not set, then only opened branches can be referenced by this field by default.</p> <ul style="list-style-type: none"> <li>Includes <p>Specifies the dataspace identifiers that can be referenced by this field.</p> <p><b>Pattern:</b> Specifies a pattern that filters dataspace identifiers. The pattern is checked against the name of the dataspace identifiers.</p> <p><b>Type:</b> Specifies the type of dataspace identifiers that can be referenced by this field. If not defined, this restriction is applied to branches.</p> <p><b>Include descendants:</b> Specifies if children or descendants of the dataspace identifiers that match the specified pattern are included in the set. If not defined, this restriction is not applied to child dataspace identifiers. If "None"</p> </li> </ul> |

then neither children nor descendants of the dataspace that match the specified pattern are included. If "All descendants" then all descendants of the dataspace that match the specified pattern are included. If "All descendant branches" then all descendant branches of the dataspace that match the specified pattern are included. If "All descendant snapshots" then all descendant snapshots of the dataspace that match the specified pattern are included. If "Child branches" then only direct branches of the dataspace that match the specified pattern are included. If the current dataspace is a version, includes the branches that are the direct children of this version; if the current dataspace is a branch, includes the branches that are the direct children of the versions which are children of this branch. If "Child snapshots" then only direct snapshots of the dataspace that match the specified pattern are included. If the current dataspace is a branch, includes the snapshots that are the direct children of this branch; if the current dataspace is a version, includes the versions that are the direct children of the branches which are children of this version.

- **Excludes**

Specifies the dataspace that cannot be referenced by this field. Excludes are ignored if no includes are defined.

**Pattern:** Specifies a pattern that filters dataspace. The pattern is checked against the name of the dataspace.

**Type:** Specifies the type of dataspace that can be referenced by this field. If not defined, this restriction is applied to branches.

**Include descendants:** Specifies if children or descendants of the dataspace that match the specified pattern are included in the set. If not defined, this restriction is not applied to child dataspace. If "None" then neither children nor descendants of the dataspace that match the specified pattern are included. If "All descendants" then all descendants of the dataspace that match the specified pattern are included. If "All descendant branches" then all descendant branches of the dataspace that match the specified pattern are included. If "All descendant snapshots" then all descendant snapshots of the dataspace that match the specified pattern are included. If "Child branches" then only direct branches of the dataspace that match the specified pattern are included. If the current dataspace is a version, includes the branches that are the direct children of this version; if the current dataspace is

a branch, includes the branches that are the direct children of the versions which are children of this branch. If "Child snapshots" then only direct snapshots of the dataspace that match the specified pattern are included. If the current dataspace is a branch, includes the snapshots that are the direct children of this branch; if the current dataspace is a version, includes the versions that are the direct children of the branches which are children of this version.

- Dataspace filter

Specifies a filter to accept or reject dataspace in the context of a dataset or a record. This filter is only used in the dedicated input component that is associated to this field. That is, this filter is not used when validating this field. A specific constraint can be used to perform specific controls on this field. A filter is defined by a Java class that implements the `DataspaceSetFilter`<sup>API</sup> interface of the Java API.

---

## Dataset set configuration

Define the datasets that can be referenced by a field of the type Dataset identifier (osd:datasetName).

- Includes

Specifies the datasets that can be referenced by this field.

**Pattern:** Specifies a pattern that filters datasets. The pattern is checked against the name of the datasets.

**Include descendants:** Specifies if children or descendants of the datasets that match the specified pattern are included in the set.

- Excludes

Specifies the datasets that cannot be referenced by this field. Excludes are ignored if no includes are defined.

**Pattern:** Specifies a pattern that filters datasets. The pattern is checked against the name of the datasets.

**Include descendants:** Specifies if children or descendants of the datasets that match the specified pattern are included in the set.

- Filter

Specifies a filter to accept or reject datasets in the context of a dataset or record. This filter is only used in the dedicated input component that is associated to this field. That is, this filter is not used when validating this field. A specific constraint can be used to perform specific controls on this field. A filter is defined by

a Java class that implements the `DatasetSetFilter`<sup>API</sup> interface of the Java API.

---

## Validation properties

Each constraint not using a specific Java class can define localized validation messages with a severity using the following properties:

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Validation</b>              | Defines a localized validation message with a user-defined severity level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Severity</b>                | Defines the severity level of the validation message. Possible values are 'Error', 'Warning', and 'Information'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Error management policy</b> | Specifies the behavior of the constraint when validation errors occur. It is possible to specify that the constraint must always remain valid after an operation (dataset update, record creation, update or deletion), or when a user submits a form. In this case, any input or operation that would violate the constraint will be rejected and the values will remain unchanged. If not specified, the constraint only blocks errors upon form submission by default, except for foreign key constraints in relational data models where errors are prevented for all operations by default. This option is only available upon static controls, exclude values, exclude segment and foreign key constraints. On foreign key constraints the error management policy does not concern filters. That is, a foreign key constraint is not blocking if a referenced record exists but does not satisfy a foreign key filter. In this case updates are not rejected and a validation error occurs. It is not possible to specify an error management policy on structural constraints that are defined in relational data models, when table history or replication is activated. That is, setting this property on fixed length, maximum length, maximum number of digits and decimal place constraints will raise an error at data model compilation because of the underlying RDBMS blocking constraints validation policy. This property is ineffective when importing archives. That is, all blocking constraints, excepted structural constraints, are always disabled when importing archives. |
| <b>Message</b>                 | Defines the message to display if the value of this field in a dataset does not comply with this constraint. If specifying a custom message, you may optionally provide localized variants.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

---

**Related concepts** [\*Properties of data model elements\*](#) [p 51]



# CHAPTER 10

---

## Toolbars

This chapter contains the following topics:

1. [Definition](#)

### 10.1 Definition

A toolbar allows to customize the buttons and menus that are displayed when viewing tables or records in a dataset. The customization of toolbars can be performed in the data model via the 'Configuration' section.

Add a toolbar from the *Toolbars* section of the navigation pane, by clicking on the arrow ▼ located to the right of [ *All elements* ], then selecting the *Create toolbar* option. Follow the creation wizard to create a toolbar. A toolbar defines the following information:

|                              |                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>                  | Toolbar's name. The name of the toolbar must be unique in the context of the data model. That is, it is not allowed to create several toolbars with the same name. |
| <b>Label and description</b> | Internationalized labels and descriptions to be displayed to end users.                                                                                            |
| <b>Default template</b>      | Allows to create a toolbar with the structure of a default toolbar.                                                                                                |
| <b>Locations</b>             | Specifies the locations where the toolbar can be used in associated datasets.                                                                                      |

#### ***Defining the structure of a toolbar***

A toolbar can define the following elements:

- [Action button](#) [p 75]
- [Menu button](#) [p 76]
- [Separator](#) [p 76]
- [Menu group](#) [p 77]

- [Action menu item](#) [p 78]
- [Sub menu item](#) [p 79]

Add one of these elements under a toolbar or to an existing element by clicking on the arrow ▼ located to the right of the existing element, and by selecting a creation option in the menu. Then, follow the creation wizard to create an element.

## Action button

This type of element allows to associate an action to a button in a toolbar. The action will be triggered when the user clicks on the associated button in one of the toolbars. A *Action button* type element defines the following information:

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Target</b>                | Defines if the service is executed in the current context or in a web component. A service executed in a web component has no access to the current selection, but it is possible to specify a selection and specific parameters.                                                                                                                                                                                                         |
| <b>Service</b>               | Defines the service that will be executed when the user clicks on the button. It is possible to select a built-in service, or a user service defined in a module or in the current data model. If the 'Web component' target is selected, the service will have to be declared as available as a web component for toolbars.<br><br><i>See also <code>WebComponentDeclarationContext.setAvailableAsToolbarAction<sup>API</sup></code></i> |
| <b>Modal size</b>            | Size of the modal window containing the web component.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Resizable</b>             | Indicates if the modal window can be resized.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>            | Service parameters. Those can be specified only if the target is set to 'Web component'.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Label and description</b> | Internationalized labels and descriptions to be displayed to end users.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Layout</b>                | Defines how this element will be displayed in datasets using the toolbar. It is possible to display: the icon only, the text only, text with the icon to the left or text with the icon to the right.                                                                                                                                                                                                                                     |
| <b>Icon</b>                  | Icon to display. It is possible to use an icon to choose from a set of suggested icons, or to refer to an icon using a URL.                                                                                                                                                                                                                                                                                                               |
| <b>Is highlighted</b>        | Indicates if the button should be highlighted by default.                                                                                                                                                                                                                                                                                                                                                                                 |

### Note

A *Action button* type element can only be created under a *toolbar* type element.

## Menu button

This type of element allows to define a menu that will be displayed when the user clicks on the associated button in a toolbar. An element of the *Menu button* type defines the following information:

|                              |                                                                                                                                                                                                       |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Label and description</b> | Internationalized labels and descriptions to be displayed to end users.                                                                                                                               |
| <b>Layout</b>                | Defines how this element will be displayed in datasets using the toolbar. It is possible to display: the icon only, the text only, text with the icon to the left or text with the icon to the right. |
| <b>Icon</b>                  | Icon to display. It is possible to use an icon to choose from a set of suggested icons, or to refer to an icon using a URL.                                                                           |
| <b>Is highlighted</b>        | Indicates if the button should be highlighted by default.                                                                                                                                             |

### Note

An element of the *Menu button* type can only be created under an element of the *toolbar* type.

## Separator

This type of element allows to insert a separator in the form of spacing between two elements of a toolbar.

### Note

An element of the *Separator* type can only be created under an element of the *toolbar* type.

## Menu group

This type of element allows to define a group of elements in a menu. An element of the *Menu group* type defines the following information:

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Label and description</b>   | Internationalized labels and descriptions to be displayed to end users.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Group type</b>              | Specifies the type of menu group to create: - 'Local' allows to create an empty fully customizable menu group. - 'Service group' allows to assign an existing service group to this menu group. - 'Menu builder' allows to assign a predefined menu content to this menu group. Once created, it is not possible to change the type of this menu group.                                                                               |
| <b>Service group name</b>      | Specifies an existing group of services to reuse. A group is declared in a module and can include other groups of services. All services contained in this group will be displayed to end users in associated datasets.                                                                                                                                                                                                               |
| <b>Menu builder name</b>       | Specifies the predefined menu content to assign to this menu group: - 'Default menu "Actions"' has the same content as the default toolbar 'Actions' menu. Standard and custom services are displayed without distinction. - 'Default menu "Actions" (with separator)' has the same menu content as above, but displays differently since standard and custom services are separated (standard services first, then custom services). |
| <b>Excluded services</b>       | Indicates the services to exclude from the group of reused services. These services will not be displayed to end users in associated datasets.                                                                                                                                                                                                                                                                                        |
| <b>Excluded service groups</b> | Indicates the groups to exclude from the group of services to reuse. Services in excluded groups will not be displayed to end users in associated datasets.                                                                                                                                                                                                                                                                           |
| <b>Filtering policy</b>        | In case of "Smart filtering", services that are configured in direct access, i.e. via an action button or an action menu item, will be removed from the automatic generation of this group.                                                                                                                                                                                                                                           |

### Note

An element of the *Menu group* type can only be created under the following elements:

- Menu button

- Sub menu item

### Action menu item

This type of element allows to associate an action to a menu item in a toolbar. The action will be triggered when the user clicks on the corresponding item in a menu. An element of the *Action menu item* type defines the following information:

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Label and description</b> | Internationalized labels and descriptions to be displayed to end users.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Target</b>                | Defines if the service is executed in the current context or in a web component. A service executed in a web component has no access to the current selection, but it is possible to specify a selection and specific parameters.                                                                                                                                                                                                         |
| <b>Service</b>               | <p>Defines the service that will be executed when the user clicks on the button. It is possible to select a built-in service, or a user service defined in a module or in the current data model. If the 'Web component' target is selected, the service will have to be declared as available as a web component for toolbars.</p> <p>See also <code>WebComponentDeclarationContext.setAvailableAsToolbarAction<sup>APT</sup></code></p> |
| <b>Modal size</b>            | Size of the modal window containing the web component.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Resizable</b>             | Indicates if the modal window can be resized.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>            | Service parameters. Those can be specified only if the target is set to 'Web component'.                                                                                                                                                                                                                                                                                                                                                  |

#### Note

An element of the *Action menu item* type can only be created under a *Menu group* type element.

### Sub menu item

This type of element allows to add a sub menu to a toolbar menu. Un *Sub menu item* defines the following information:

| Label and description | Internationalized labels and descriptions to be displayed to end users. |
|-----------------------|-------------------------------------------------------------------------|
|-----------------------|-------------------------------------------------------------------------|

**Note**

An element of the *Sub menu item* type can only be created under an element of the Menu group type.

### Deleting elements

All the elements of a toolbar can be deleted from it by using the arrow ▼ located to the right of the element to be deleted.

If an element containing other elements is deleted, then the deletion is recursively performed on all elements located below the deleted element.

### Duplicating existing elements

To duplicate an element, click on the arrow ▼ located to the right of the element to duplicate. Specify the name and properties of the duplicated element. All the source element properties are duplicated.

The duplicated element is added on the same level than the original element, in the final position. When an element containing other elements is duplicated, all the sub-elements are duplicated with their properties.

### Moving elements

In order to move an element, click on the arrow ▼ and select the moving option to be used.

### Associate with existing tables

To associate a toolbar with existing tables, click on the arrow ▼ located to the right of the toolbar and select the option *Associate to tables*. This service allows to set the toolbar has the default toolbar of several tables in one shot. To do so, specify the target locations of the toolbar and select the tables or complex data types, that define table properties, to be associated with the toolbar.

## Exporting the toolbars

It is possible to export the toolbars defined in the model into an XML document. To do so, select the *XML export* option available in the *Actions* menu of the 'Toolbars' section. Follow the wizard to export the toolbars.

### Note

A selection of toolbars can be exported by selecting in the 'Toolbars' section the toolbars to be exported and then by selecting the *XML export* option available in the *Actions* menu. The toolbars can also be exported by using the data model export service. It can be found in the [Data model 'Actions'](#) [p 33] menu in the navigation pane.

See also [XML Schema Document \(XSD\) import and export](#) [p 81]

## Importing toolbars

It is possible to import existing toolbars from an XML document. To do so, select the *XML import* option available in the *Actions* menu of the 'Toolbars' section. Then follow the wizard to import the toolbars.

### Note

The toolbars can also be imported by using the data model import service accessible via the [Data model 'Actions'](#) [p 33] menu in the navigation pane.

See also [XML Schema Document \(XSD\) import and export](#) [p 81]

See also [Use of toolbars](#) [p 61]



## CHAPTER 11

# Working with an existing data model

Once your data model has been created, you can perform a number of actions that are available from the [data model 'Actions'](#) [p 33] menu in the workspace.

This chapter contains the following topics:

1. [Validating a data model](#)
2. [XML Schema Document \(XSD\) import and export](#)
3. [Duplicating a data model](#)
4. [Deleting a data model](#)

## 11.1 Validating a data model

To validate a data model at any time, select **Actions > Validate** from the navigation pane. The generated report provides the results of the validation. From the validation report, you have the option to update the reported validation status by clicking the **Revalidate** button, or to click the **Reinitialize validation report** button to clear the current validation report associated with the data model in order to be able to rerun a full validation from scratch.

**Note**

The validation process checks basic data model integrity, but more complex checks are only performed at publication time. More messages may be reported when you try to publish your data model.

See [Validation](#) [p 298] for detailed information on incremental data validation.

## 11.2 XML Schema Document (XSD) import and export

TIBCO EBX® includes built-in data model services to import from and export to XML Schema Document (XSD) files. XSD imports and exports can be performed from the [data model 'Actions'](#) [p 33] menu of the target data model in the navigation pane. An XSD import or export is always performed in the context of a single data model. That is, during imports, the structure of the target data model is completely replaced with the content of the imported XSD. Similarly, during exports, the entire data model is included in the XSD file.

When importing an XSD file, the file must be well-formed and must comply with EBX® validation rules. If this document declares resources that are located in a module, the module must also be declared in the configuration of the data model. If the module has not been declared, you will not

be able to import the XSD file. See [Data model properties](#) [p 40] for more information on declaring modules.

To perform an import select 'Import XSD' from the [data model 'Actions'](#) [p 33] menu of the data model into which you are importing.

You can import an XML Schema Document (XSD) from the local file system. To do so, click the 'From a local document' button in the import wizard and follow the next step:

- **Document name:** path on the local file system of the XSD file to import.

You can also import a data model in an XSD that is packaged in a module. The import of a data model in XSD format from a module uses the following properties:

|                    |                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Module</b>      | Module in which the data model is packaged.                                                                                                                        |
| <b>Module path</b> | Path to the module containing the data model.                                                                                                                      |
| <b>Source path</b> | Path to Java source used to configure business objects and rules.<br><br>This property is required if the data model being imported defines programmatic elements. |
| <b>Model</b>       | The data model in the module to import.                                                                                                                            |

#### Note

Imported XSD files must be encoded in 'UTF-8'. Exported XSD files are always encoded in 'UTF-8'.

## 11.3 Duplicating a data model

To duplicate a data model, select 'Duplicate' from the [data model 'Actions'](#) [p 33] menu for that data model. You must give the new data model a name that is unique in the repository.

## 11.4 Deleting a data model

To delete a data model, select 'Delete' from the [data model 'Actions'](#) [p 33] menu for that data model. When you delete a data model, all of its existing publications will remain and continue to be accessible to their associated datasets. If you recreate a new data model with the same name as one that was previously deleted, the new data model will be reassociated with all the existing publications in the repository. At publication time of the new data model, you will have the opportunity to confirm the replacement of an existing publication.

#### Note

Only an administrator can clean up the publications of deleted data models in the 'Administration' area.

See [Publishing data models](#) [p 83] for more information on the publication process.

## CHAPTER 12

# Publishing a data model

This chapter contains the following topics:

1. [About publications](#)
2. [Publication modes](#)
3. [Embedded publication mode](#)

## 12.1 About publications

Each dataset based on an **embedded data model** in the TIBCO EBX® repository is associated with a publication of a data model, rather than directly to the data model itself. The first time you publish a data model using the **Publish** button in the navigation pane, a publication is created. Once the publication exists, datasets can be created based upon it.

**Note**

The **Publish** button is only displayed to users who have permission to publish the data model. See [Data model permissions](#) [p 39] for more information.

As datasets are based on publications, any modifications you make to a data model will only take effect on existing datasets when you republish to the publication associated with those datasets. When you republish a data model to an existing publication, all existing datasets associated with that particular publication are updated.

## 12.2 Publication modes

You can publish a data model using either 'Embedded' mode or 'In module' mode. The 'Embedded' publication mode generates a publication that is managed and persisted within the EBX® repository and thus has associated versioning and rollback functionality. The 'In module' publication mode creates an XML Schema Document contained in a module that is not managed or versioned within the repository.

Depending on the configuration of the data model, EBX® automatically determines the publication process to use when you click the **Publish** button in the navigation pane. When a data model specifies the publication mode 'In module' and provides a target XSD to be generated, the publication process generates an XSD file contained in the module defined in the configuration.

## 12.3 Embedded publication mode

The first time you publish a given embedded data model, a new publication with the same name as your data model is automatically created in the repository. If more than one publication has already been created for this model, you will need to select a target publication for this process.

See [Viewing and creating publications](#) [p 84] for more information on the use of different publications.

During the publication process, you have the opportunity to review the structural differences being introduced by the current publication in a side-by-side comparison view, if the data model has previously been published.

The publication process also offers the option to create a read-only snapshot of the current state of the data model for reference purposes. This snapshot can be useful if the data model ever needs to be rolled back to the current state after other modifications have been made.

### Note

Snapshots, which are static archives of the state of the data model, must not be confused with data model *versions*, which act instead as parallel evolving branches of the data model. See [Versioning embedded data models](#) [p 85] for more information on data model versions.

### *Viewing and creating publications*

To access the publications that exist for the current data model, select 'Manage publications' from its [data model 'Actions'](#) [p 33] menu in the navigation pane. From there, you can view the details of the publications and create new publications.

In certain cases, it may be necessary to employ several publications of the same data model, in order to allow datasets to be based on different states of that data model. Multiple publications must be handled carefully, as users will be asked to select an available publications to target when publishing if more than one exists. The action to create a new publication is only available to users who belong to the 'Administrator' role.

To create a new publication, select 'Manage publications' from the [data model 'Actions'](#) [p 33] menu of the data model in the navigation pane, then click the **Create publication** button. The name you give to the publication must unique in the repository.

## CHAPTER 13

# Versioning an embedded data model

This chapter contains the following topics:

1. [About versions](#)
2. [Accessing versions](#)
3. [Working with versions](#)
4. [Known limitations on data model versioning](#)

## 13.1 About versions

You can create *versions* for data models that evolve in parallel. Versions are not to be confused with data model snapshots, which are taken at publication time and kept strictly for historical read-only reference.

## 13.2 Accessing versions

To see the existing versions of your data model, select 'Manage versions' from the [data model](#) **'Actions'** [p 33] menu of the data model.

The existing versions are represented in a tree format according to their parent-child relationships. Every data model has a root version by default, with the same name as the data model.

## 13.3 Working with versions

In the workspace, using the down arrow ▼ menu next to each version, you can perform the following actions:

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Access data model version</b> | Go to the corresponding version of the data model.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Create version</b>            | Creates a new version based on the contents of the selected version. The new version is added as a child of the selected version, though its contents bear no relation to those of its parent version after creation.                                                                                                                                                                               |
| <b>Set as default version</b>    | Sets the selected version as the default version opened when users access the data model.                                                                                                                                                                                                                                                                                                           |
| <b>Export archive</b>            | Exports the selected data model version to an archive containing the version's content, including its permissions and information. The exported archive is located in the archives directory, which is accessible to repository administrators. Exporting to an existing archive name will overwrite the existing file.<br><br>See <a href="#">Archives directory</a> [p 378] for more information. |
| <b>Import archive</b>            | Imports the content of an archive into the selected version. The archive to import must contain a data model with the same name as the data model associated with the version.                                                                                                                                                                                                                      |

A version can be deleted by clicking the **X** button to the right of its entry. A version cannot be deleted if it is linked to a publication or if it has child versions. The root version of a data model also cannot be deleted.

Two versions of the same data model can be compared in the workspace by selecting their checkboxes, then selecting **Actions > Compare selected versions**. The side-by-side comparison shows structural differences between the version of the data model, with the older version on the left and the newer version on the right.

## 13.4 Known limitations on data model versioning

- It is not possible to merge two versions of a data model.
- The comparison interface does not display updates on fields, only additions and deletions.
- Versioning of data models packaged in modules is not supported.
- Resources packaged in a module that are used by an embedded data model are not versioned when a version is created. That is, only the reference of the resources are saved during the creation of a version, and it is the responsibility of developers to ensure that the content of the referenced resources are compatible with any versions that may be using them.

---

# Dataspaces

---

## CHAPTER 14

---

# Introduction to dataspaces

This chapter contains the following topics:

1. [Overview](#)
2. [Using the Dataspaces area user interface](#)

## 14.1 Overview


### ***What is a dataspace?***

The life cycle of data can be complex. It may be necessary to manage a current version of data while working on several concurrent updates that will be integrated in the future, including keeping a trace of various states along the way. In TIBCO EBX®, this is made possible through the use of dataspaces and snapshots.

A dataspace is a container that isolates different versions of datasets and organizes them. A dataspace can be branched by creating a child dataspace, which is automatically initialized with the state of its parent. Thus, modifications can be made in isolation in the child dataspace without impacting its parent or any other dataspaces. Once modifications in a child dataspace are complete, that dataspace can be compared with and merged back into the parent dataspace.



Snapshots, which are static, read-only captures of the state of a dataspace at a given point in time, can be taken for reference purposes. Snapshots can be used to revert the content of a dataspace later, if needed.



### ***Basic concepts related to dataspaces***

A basic understanding of the following terms is beneficial when working with dataspaces:

- [dataspace](#) [p 26]
- [snapshot](#) [p 26]
- [dataset](#) [p 24]
- [dataspace merge](#) [p 26]
- [reference dataspace](#) [p 26]


## **14.2 Using the Dataspaces area user interface**

Dataspaces can be created, accessed and modified in the **Dataspaces** area.

#### **Note**

This area is available only to authorized users in the 'Advanced perspective'.

The navigation pane displays all existing dataspaces, while the workspace displays information about the selected datasource and lists its snapshots.



See also

[Creating a datasource](#) [p 91]

[Snapshots](#) [p 101]

Related concepts [Datasets](#) [p 106]

## CHAPTER 15

# Creating a dataspace

This chapter contains the following topics:

1. [Overview](#)
2. [Properties](#)
3. [Relational mode](#)

## 15.1 Overview

By default, dataspaces in TIBCO EBX® are in *semantic mode*. This mode offers full-featured data life cycle management.

To create a new dataspace in the default semantic mode, select an existing dataspace on which to base it, then click the **Create a dataspace** button in the workspace.

### Note

This area is available only to authorized users in the 'Advanced perspective'.

The new dataspace will be a child dataspace of the one from which it was created. It will be initialized with all the content of the parent at the time of creation, and an initial snapshot will be taken of this state.

Aside from the reference dataspace, which is the root of all semantic dataspaces in the repository, semantic dataspaces are always a child of another dataspace.

See also [Relational mode](#) [p 92]

## 15.2 Properties

The following information is required at the creation of a new dataspace:

|                        |                                                                                                                                                                        |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Identifier</b>      | Unique identifier for the dataspace.                                                                                                                                   |
| <b>Owner</b>           | Owner of the dataspace, who is, by default, allowed to modify its information and permissions. The owner does not necessarily have to be the creator of the dataspace. |
| <b>Label</b>           | Localized label and description associated with the dataspace.                                                                                                         |
| <b>Relational mode</b> | Whether or not this dataspace is in relational mode. This option only exists when creating a new dataspace from the reference dataspace.                               |

## 15.3 Relational mode

Dataspaces in relational mode can only be created from the reference dataspace. They offer limited functionality compared to dataspaces in semantic mode. For instance, dataspaces in relational mode do not handle snapshots or support child dataspaces.

Relational mode indicates that the tables of data in this dataspace are stored directly in an RDBMS.

See also [Relational mode](#) [p 243]

## CHAPTER 16

---

# Working with existing dataspaces

This chapter contains the following topics:

1. [Dataspace information](#)
2. [Dataspace permissions](#)
3. [Merging a dataspace](#)
4. [Comparing a dataspace](#)
5. [Validating a dataspace](#)
6. [Dataspace archives](#)
7. [Closing a dataspace](#)

## 16.1 Dataspace information

Certain properties associated with a dataspace can be modified by selecting **Actions > Information** from the navigation panel in the Dataspaces area.

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Documentation</b>               | Localized labels and descriptions associated with the dataspace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Loading strategy</b>            | <p><i>Only administrators can modify this setting.</i></p> <ul style="list-style-type: none"> <li>• <b>On demand loading and unloading:</b> The main advantage of this strategy is the ability to free memory when needed. Its disadvantage is the performance cost associated with a resource being accessed for the first time since server startup, or accessed after having been unloaded. This is the default mode.</li> <li>• <b>Forced loading:</b> This mode is recommended for dataspaces and snapshots used heavily or demanding in terms of response time.</li> <li>• <b>Forced loading and prevalidation:</b> This mode is recommended for dataspaces and snapshots used heavily or demanding in terms of response time, and where the validation process can be time-intensive.</li> </ul> <p><b>Note:</b> Whenever the loading strategy is changed, you must restart the server for the new setting to take effect.</p> |
| <b>Child merge policy</b>          | <p>This merge policy only applies to user-initiated merge processes; it does not apply to programmatic merges, for example, those performed by workflow script tasks.</p> <p>The available merge policies are:</p> <ul style="list-style-type: none"> <li>• <b>Allows validation errors in result:</b> Child dataspaces can be merged regardless of the validation result. This is the default policy.</li> <li>• <b>Pre-validating merge:</b> A child dataspace can only be merged if the result would be valid.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Current Owner</b>               | Owner of the dataspace, who is, by default, allowed to modify its information and permissions. The owner does not necessarily have to be the creator of the dataspace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Child dataspace sort policy</b> | Defines the display order of child dataspaces in dataspace trees. If not defined, the policy of the parent dataspace is applied. Default is 'by label'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

|                           |                                                                                                                                                                                                                |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Change owner</b>       | Whether the current owner of the dataspace is allowed to give ownership to another profile by modifying its 'Current owner' property. If the value is 'Forbidden', only an administrator can modify the owner. |
| <b>Change permissions</b> | Whether the current owner of the dataspace is allowed to modify its permissions. If the value is 'Forbidden', only an administrator can modify the permissions of the dataspace.                               |

## 16.2 Dataspace permissions

### *General permissions*

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Dataspace id</b>       | The dataspace to which the permissions will apply.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Profile selection</b>  | The profile to which the rule applies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Restriction policy</b> | Whether these permissions restrict the permissions assigned to a given user through policies defined for other profiles.<br>See <a href="#">Restriction policy</a> [p 281].                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Dataspace access</b>   | <p>The global access permission on the dataspace.</p> <p><b>Read-only</b></p> <ul style="list-style-type: none"> <li>• Can see the dataspace and its snapshots, as well as child dataspace, according to their permissions.</li> <li>• Can see the contents of the dataspace depending on their permissions; cannot make modifications.</li> </ul> <p><b>Write</b></p> <ul style="list-style-type: none"> <li>• Can see the dataspace and its snapshots, as well as child dataspace, according to their permissions.</li> <li>• Can modify the contents of the dataspace depending on their permissions.</li> </ul> <p><b>Hidden</b></p> <ul style="list-style-type: none"> <li>• Cannot see the dataspace nor its snapshots directly.</li> <li>• From a child dataspace, the current dataspace can be seen but not selected.</li> <li>• Cannot access the contents of the dataspace.</li> <li>• Cannot perform any actions on the dataspace.</li> </ul> |

## Allowable actions

Users can be allowed to perform the following actions:

|                                                    |                                                                                                           |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Create a child dataspace</b>                    | Whether the profile can create child dataspace.                                                           |
| <b>Create a snapshot</b>                           | Whether the profile can create snapshots from the dataspace.                                              |
| <b>Initiate merge</b>                              | Whether the profile can merge the dataspace with its parent.                                              |
| <b>Export archive</b>                              | Whether the profile can perform exports.                                                                  |
| <b>Import archive</b>                              | Whether the profile can perform imports.                                                                  |
| <b>Close dataspace</b>                             | Whether the profile can close the dataspace.                                                              |
| <b>Close snapshot</b>                              | Whether the profile can close snapshots of the dataspace.                                                 |
| <b>Rights on services</b>                          | Specifies the access permissions for services.                                                            |
| <b>Permissions of child dataspace when created</b> | Specifies the default access permissions for child dataspace that are created from the current dataspace. |

## 16.3 Merging a dataspace

When the work in a given dataspace is complete, you can perform a one-way merge of the dataspace back into the dataspace from which it was created. The merge process is as follows:

- Both the parent and child dataspace are locked to all users, except the user who initiated the merge and administrator users. These locks remain for the duration of the merge operation. When locked, the contents of a dataspace can be read, but they cannot be modified in any way.  
**Note:** This restriction on the parent dataspace means that, in addition to blocking direct modifications, other child dataspace cannot be merged until the merge in progress is finished.
- Changes that were made in the child dataspace since its creation are integrated into its parent dataspace.
- The child dataspace is closed.
- The parent dataspace is unlocked.

### Initiating a merge

To merge a dataspace into its parent dataspace:

- Select that dataspace in the navigation pane of the Dataspace area.
- In the workspace, select **Merge dataspace** from the **Actions** menu.



## Reviewing and accepting changes

After initiating a dataspace merge, you must review the changes that have been made in the child (source) dataspace since its creation, to decide which of those changes to apply to the parent (target) dataspace.

### Note

This change set review and acceptance stage is bypassed when performing merges using data services or programmatically. For automated merges, all changes in the child dataspace override the data in the parent dataspace.

The change acceptance process uses a side-by-side comparison interface that recapitulates the changes that require review. Two *change set* columns are obtained by taking the relevant changes from the following dataspace state comparisons:

- The current child dataspace compared to its initial snapshot.
- The parent dataspace compared to the initial snapshot of the child dataspace.

By default, all detected changes are selected to be merged. You may deselect any changes that you want to omit from the merge. You can view the changes relevant to different scopes in your data model by selecting elements in the navigation pane.

In order to detect conflicts, the merge involves the current dataspace, its initial snapshot and the parent dataspace, because data is likely to be modified both in the current dataspace and its parent.

The merge process also handles modifications to permissions on tables in the dataspace. As with other changes, access control changes must be reviewed for inclusion in the merge.

When you have decided which changes to merge for a given scope, you must click the button **Mark difference(s) as reviewed** to indicate that you have reviewed all the changes in that scope. All changes must be reviewed in order to proceed with the merge.

## Types of modifications

The merge process considers the following changes as modifications to be reviewed:

- Record and dataset creations
- Any changes to existing data
- Record, dataset, or value deletions
- Any changes to table permissions

## Types of conflicts

This review interface also shows conflicts that have been detected. Conflicts may arise when the same scope contains modifications in both the source and target dataspace.

Conflicts are categorized as follows:

- A record or a dataset creation conflict
- An entity modification conflict
- A record or dataset deletion conflict
- All other conflicts

## Finalizing a merge

Once you have reviewed all changes and decided which to include in the merge result, click on the **Merge >>** button in the navigation pane.

Depending on the child merge policy that is configured on the parent dataspace in your merge, the subsequent process may differ. By default, merges are finalized even if the result would contain validation errors. The administrator of the parent dataspace in your merge can set its child merge policy so that merges of its children are only finalized if the result would not contain any validation errors.

If, however, the administrator of the parent dataspace has set its child merge policy to 'Pre-validating merge', a dedicated dataspace is first created to hold the result of the merge. When the result is valid, this dedicated dataspace containing the merge result is automatically merged into the parent dataspace, and no further action is required.

In the case where validation errors are detected in the dedicated merge dataspace, you only have access to the original parent dataspace and the dataspace containing the merge result, named "[merge] < name of child dataspace >". The following options are available to you from the **Actions > Merge in progress** menu in the workspace:

- **Cancel**, which abandons the merge and recuperates the child dataspace in its pre-merge state.
- **Continue**, which you can use to reattempt the merge after you have made the necessary corrections to the dedicated merge dataspace.

## Setting the child merge policy of a dataspace

As the administrator of a dataspace, you can block the finalization of merges of its child dataspace through the user interface when the merges would result in a dataspace with validation errors. To do so, select **Actions > Information** from the workspace of the parent dataspace. On the dataspace's information page, set the **Child merge policy** to **Pre-validating merge**. This policy will then be applied to the merges of all child dataspace into this parent dataspace.

### Note

When the merge is performed through a Web Component, the behavior of the child merge policy is the same as described; the policy defined in the parent dataspace is automatically applied when merging a child dataspace. However, this setting is ignored during programmatic merge, which includes script tasks in data workflows.

See also [Child merge policy](#) [p 98]

## Abandoning a merge

Merges are performed in the context of a user session, and must be completed in a single operation. If you decide not to proceed with a merge that you have initiated, you can click the **Cancel** button to abandon the operation.

If you navigate to another page after initiating a merge, the merge will be abandoned, but the locks on the parent and child dataspace will remain until you unlock them in the Dataspace area.

You may unlock a dataspace by selecting it in the navigation pane, and clicking the **Unlock** button in the workspace. Performing the unlock from the child dataspace unlocks both the child and parent dataspace. Performing the unlock from the parent dataspace only unlocks the parent dataspace, thus you need to unlock the child dataspace separately.

## 16.4 Comparing a dataspace

You can compare the contents of a dataspace to those of another dataspace or snapshot in the repository. To perform a comparison, select the dataspace in the navigation pane, then select **Actions > Compare** from the workspace.

The comparison wizard prompts you to select the dataspace or snapshot with which to compare the current dataspace.

For a faster comparison that ignores fields with inherited and computed values, select the filter 'Persisted values only'.

See also [Compare contents](#) [p 209]

## 16.5 Validating a dataspace

To perform a global validation of the contents of a dataspace, select that dataspace in the navigation panel, then select **Actions > Validate** in the workspace.

### Note

This service is only available in the user interface if you have permission to validate every dataset contained in the current dataspace.

## 16.6 Dataspace archives

The content of a dataspace can be exported to an archive or imported from an archive.

### Exporting

To export a dataspace to an archive, select that dataspace in the navigation panel, then select **Actions > Export** in the workspace. Once exported, the archive file is saved to the file system of the server, where only an administrator can retrieve the file.

### Note

See [Archives directory](#) [p 378] in the Administration Guide for more information.

In order to export an archive, the following information must be specified:

|                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name of the archive to create</b> | The name of the exported archive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Export policy</b>                 | <p>Required.</p> <p>The default export policy is 'The whole content of the dataspace', which exports all selected data to the archive.</p> <p>It may be useful to export only the differences between the dataspace and its initial snapshot using a change set. There are two different export options that include a change set: 'The updates with their whole content' and 'The updates only'. The first option exports all current data and a change set containing differences between the current state and the initial snapshot. The second option only exports the change set. Both options lead to a comparison page, where you can select the differences to include in this change set. Differences are detected at the table level.</p> |
| <b>Datasets to export</b>            | The datasets to export from this dataspace. For each dataset, you can export its data values, permissions, and/or information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Importing

To import content into a dataspace from an archive, select that dataspace in the navigation panel, then select **Actions > Import** in the workspace.

If the selected archive does not include a change set, the current state of the dataspace will be replaced with the content of the archive.

If the selected archive includes the whole content as well as a change set, you can choose to apply the change set in order to merge the change set differences with the current state. Applying the change set leads to a comparison screen, where you can then select the change set differences to merge.

If the selected archive only includes a change set, you can select the change set differences to merge on a comparison screen.

## 16.7 Closing a dataspace

If a dataspace is no longer needed, it can be closed. Once it is closed, a dataspace no longer appears in the **Dataspaces** area of the user interface, nor can it be accessed.

An administrator can reopen a closed dataspace as long as it has not been cleaned from the repository.

To close a dataspace, select **Actions > Close this dataspace**.

See also [Closing unused dataspace and snapshots](#) [p 380]

# CHAPTER 17

---

## Snapshots

This chapter contains the following topics:

- 1. [Overview of snapshots](#)
- 2. [Creating a snapshot](#)
- 3. [Viewing snapshot contents](#)
- 4. [Snapshot information](#)
- 5. [Comparing a snapshot](#)
- 6. [Validating a snapshot](#)
- 7. [Export](#)
- 8. [Closing a snapshot](#)

### 17.1 Overview of snapshots

A snapshot is a read-only copy of a dataspace. Snapshots exist as a record of the state and contents of a dataspace at a given point in time.

See also [Snapshot](#) [p 26]

### 17.2 Creating a snapshot

A snapshot can be created from a dataspace by selecting that dataspace in the navigation pane of the Dataspaces area, then selecting **Actions > Create a Snapshot** in the workspace.

The following information is required:

|                   |                                                                 |
|-------------------|-----------------------------------------------------------------|
| <b>Identifier</b> | Unique identifier for the snapshot.                             |
| <b>Label</b>      | Localized labels and descriptions associated with the snapshot. |

## 17.3 Viewing snapshot contents

To view the contents of a snapshot, select the snapshot, then select **Actions > View datasets** from the workspace.

## 17.4 Snapshot information

You can modify the information associated with a snapshot by selecting **Actions > Information**.

|                         |                                                                                                                                                                                                               |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Documentation</b>    | Localized labels and descriptions associated with the snapshot.                                                                                                                                               |
| <b>Loading strategy</b> | <i>Only administrators can modify this setting. See <a href="#">Loading strategy</a> [p 94].</i>                                                                                                              |
| <b>Current Owner</b>    | Owner of the snapshot, who is, by default, allowed to modify its information and permissions. The owner does not necessarily have to be the creator of the snapshot.                                          |
| <b>Change owner</b>     | Whether the current owner of the snapshot is allowed to give ownership to another profile by modifying its 'Current owner' property. If the value is 'Forbidden', only an administrator can modify the owner. |

## 17.5 Comparing a snapshot

You can compare the contents of a snapshot to those of another snapshot or dataspace in the repository. To perform a comparison, select the snapshot, then select **Actions > Compare** from the workspace.

The comparison wizard prompts you to select the dataspace or snapshot with which to compare the current snapshot.

For a faster comparison that ignores fields with inherited and computed values, select the filter 'Persisted values only'.

See also [Compare contents](#) [p 209]

## 17.6 Validating a snapshot

To perform a global validation of the contents of a snapshot, select **Actions > Validate** in the workspace.

### Note

In order to use this service, you must have permission to validate every dataset contained in the snapshot.

# 17.7 Export

To export a snapshot to an archive, open that snapshot, then select **Actions > Export** in the workspace. Once exported, only an administrator can retrieve the archive.

**Note**

See [Archives directory](#) [p 378] in the Administration Guide for more information.

In order to export an archive, the following information must be specified:

|                                      |                                                                                                                                              |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name of the archive to create</b> | The name of the exported archive.                                                                                                            |
| <b>Datasets to export</b>            | The datasets to export from this snapshot. For each dataset, you can choose whether to export its data values, permissions, and information. |

# 17.8 Closing a snapshot

If a snapshot is no longer needed, it can be closed. Once it is closed, a snapshot no longer appears under its associated dataspace in the Dataspaces area, nor can it be accessed.

An administrator can reopen a closed dataspace as long as it has not been cleaned from the repository.

To close a snapshot, select **Actions > Close this snapshot**.

See also [Closing unused dataspace and snapshots](#) [p 380]





---

# Datasets

---

## CHAPTER 18

# Introduction to datasets

This chapter contains the following topics:

1. [Overview](#)
2. [Using the Data user interface](#)

## 18.1 Overview

### ***What is a dataset?***

A dataset is a container for data that is based on the structural definition provided by its underlying data model. When a data model has been published, it is possible to create datasets based on its definition. If that data model is later modified and republished, all its associated datasets are automatically updated to match.

In a dataset, you can consult actual data values and work with them. The views applied to tables allow representing data in a way that is most suitable to the nature of the data and how it needs to be accessed. Searches and filters can also be used to narrow down and find data.

Different permissions can also be accorded to different roles to control access at the dataset level. Thus, using customized permissions, it would be possible to allow certain users to view and modify a piece of data, while hiding it from others.


### ***Basic concepts related to datasets***

A basic understanding of the following terms is beneficial when working with datasets:

- [dataspace](#) [p 26]
- [dataset](#) [p 24]
- [record](#) [p 24]
- [field](#) [p 23]
- [primary key](#) [p 23]
- [foreign key](#) [p 23]
- [table \(in dataset\)](#) [p 24]
- [group](#) [p 23]

## 18.2 Using the Data user interface

Datasets can be created, accessed and modified in the **Data** area using the [Advanced perspective](#) [p 15] or from a specifically configured perspective. Only authorized users can access these interfaces.



Select or create a dataset using the 'Select dataset' menu in the navigation pane. The data structure of the dataset is then displayed in the navigation pane, while record forms and table views are displayed in the workspace.

When viewing a table of the dataset in the workspace, the button  displays searches and filters that can be applied to narrow down the records that are displayed.

Operations at the dataset level are located in the **Actions** menu in the navigation pane (services are available at the bottom of the list).

### See also

[Creating a dataset](#) [p 109]

[Searching and filtering data](#) [p 112]

[Working with records in the user interface](#) [p 119]

[Inheritance](#) [p 25]

### Related concepts

[Data model](#) [p 32]

[Dataspace](#) [p 88]



## CHAPTER 19

# Creating a dataset

This chapter contains the following topics:

1. [Creating a root dataset](#)
2. [Creating an inheriting child dataset](#)

## 19.1 Creating a root dataset

To create a new root dataset, that is, one that does not inherit from a parent dataset, select the '[Select dataset](#) [p 107]' ▾ menu in the navigation pane, click the '**Create a dataset**' button in the pop-up, and follow through the wizard.

**Note**

This area is available only to authorized users in the 'Advanced perspective' or from a specifically configured perspective.

The wizard allows you to select one of three data model packaging modes on which to base the new dataset: packaged, embedded, or external.

- A *packaged data model* is a data model that is located within a module, which is a web application.
- An *embedded data model* is a data model that is managed entirely within the TIBCO EBX® repository.
- An *external data model* is one that is stored outside of the repository and is referenced using its URI.

After locating the data model on which to base your dataset, you must provide a unique name, without spaces or special characters. Optionally, you may provide localized labels for the dataset, which will be displayed to users in the user interface depending on their language preferences.

**Attention**

Table contents are not copied when duplicating a dataset.

## 19.2 Creating an inheriting child dataset

The inheritance mechanism allows datasets to have parent-child relationships, through which default values are inherited from ancestors by descendants. In order to be able to create child datasets, dataset inheritance must be enabled in the underlying data model.

To create a child dataset, select the '[Select dataset](#) [p 107]' ▼ menu in the navigation pane, then click the + button next to the desired parent dataset.

As the dataset will automatically be based on the same data model as the parent dataset, the only information that you need to provide is a unique name, and optionally, localized labels.

See also [Dataset inheritance](#) [p 125]

## CHAPTER 20

# Viewing table data

TIBCO EBX® offers a customization mechanism for tables via the 'Views' feature. A view allows specifying which columns should be displayed as well as the display order. Views can be managed by profile thanks to the [recommended views](#) [p 116] concept.

This chapter contains the following topics:

1. ['View' menu](#)
2. [Sorting data](#)
3. [Searching and filtering data](#)
4. [Views](#)
5. [Views management](#)
6. [Grid edit](#)
7. [History](#)

## 20.1 'View' menu

The 'View' drop-down menu allows accessing all available views and management features.

Views are managed in a dedicated sub-menu: ['Manage views'](#) [p 117].

Views can also be grouped. An administrator has to define groups beforehand in 'Views configuration' under the 'Groups of views' table. The end-user can then set a view as belonging to a group through the field 'View group' upon creation or modification of the view. See ['View description'](#) [p 114] for more information.

## 20.2 Sorting data

To simply sort a single column in a table, click on the column title. The first click will sort by ascending order and a second click will reverse the sorting.

Note that the default order is by ascending primary key.

For more advanced sorting, custom sort criteria allow specifying the display order of records.

To define custom sort criteria, click on the 'Select and sort' button in the workspace.

Each sort criterion is defined by a column name and a sort direction, that is, ascending or descending. Use the 'Move left' or 'Move right' arrows to add or remove a criterion from the 'Sorted' table. When

a criterion is highlighted, you can set its sort direction by clicking on the 'ASC' or 'DESC' button to the right.

To change the priority of a sort criterion, highlight it in the list, then use the up and down arrow buttons to move it.

To remove a custom sort order that is currently applied, select *View > Reset*.

## 20.3 Searching and filtering data

The feature for searching and filtering records is accessible via the icon  in the workspace.

When criteria are defined for a search or filter, a checkbox appears in the title bar of the pane to apply the filter. When unchecked, the search or filter is not applied.

### Note

Applying a view resets and removes all currently applied searches and filters.

### Search

In simple mode, the 'Search' tool allows adding type-contextual search criteria to one or more fields. Operators relevant to the type of a given field are proposed when adding a criterion.

By enabling the advanced mode, it is possible to build sub-blocks containing criteria for more complex logical operations to be involved in the search results computation.

### Note

In advanced mode, the criteria with operators "matches" or "matches (case sensitive)" follow the standard regular expression syntax from Java.

See also [Regex pattern](#)

### Text search

The text search is intended for plain-text searches on one or more fields. The text search does not take into account the types of the fields being searched for.

- If the entered text contains one or more words without wildcard characters (\* or ?), matching fields must contain all specified words. Words between quotes, for example "aa bb", are considered to be a single word.
- Standard wildcard characters are available: \* (any text) or ? (any character). For performance reasons, only one of these characters can be entered in each search.
- Wildcard characters themselves can be searched for by escaping them with the character '\'. For example '\\*' will search for the asterisk character.

Examples:

- aa bb: field contains 'aa' and 'bb'.
- aa "bb cc": field contains 'aa' and 'bb cc'.
- aa\*: field label starts with 'aa'.
- \*bb: field label ends with 'bb'.
- aa\*bb: field label starts with 'aa' and ends with 'bb'.



- aa?: field label starts with 'aa' and is 3 chars long.
- ?bb: field label ends with 'bb' and is 3 chars long.
- aa?bb: field label starts with 'aa' and ends with 'bb' and is 5 chars long.
- aa\\*bb: field contains 'aa\*bb' as is.

For large tables, it is recommended to select only one field, and for cases where the field type is not a string, to try to match the format type. For example:

- boolean: Yes, No
- date: 01/01/2000
- numeric: 100000 or 100,000
- enumerated value: Red, Blue...

The text search can be made case sensitive, that is distinguishing between upper and lower case, by checking the 'Case sensitive' checkbox.

### ***Validation messages filter***

The validation messages filter allows viewing records according to their status as of the last validation performed. Available levels are: 'Errors', 'Warnings', or 'Information'.

#### **Note**

This filter only applies to records of the table that have been validated at least once by selecting *Actions > Validate* at the table level from the workspace, or at the dataset level from the navigation pane.

### ***Custom table searches***

Additional custom filters can be specified for each table in the data model.

See also [Specifying UI filters on a table](#) [p 188]

## **20.4 Views**

It is possible to customize the display of tables in EBX® according to the target user. There are two types of views: [tabular](#) [p 114] and [hierarchical](#) [p 115].

A view can be created by selecting *View > Create a new view* in the workspace. To apply a view, select it in *View > name of the view*.

Two types of views can be created:

- 'Simple tabular view': A table view to sort and filter the displayed records.
- 'Hierarchical view': A tree view that links data in different tables based on their relationships.

## View description

When creating or updating a view, the first page allows specifying general information related to the view.

|                      |                                                                                                                                                                       |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Documentation</b> | Localized label and description associated with the view.                                                                                                             |
| <b>Owner</b>         | Name of the owner of the view. This user can manage and modify it. (Only available for administrators and dataset owners)                                             |
| <b>Share with</b>    | Other profiles allowed to use this view from the 'View' menu.<br><br><div><b>Note</b><br/>Requires a permission, see <a href="#">Views permissions</a> [p 399].</div> |
| <b>View mode</b>     | Simple tabular view or hierarchical view.                                                                                                                             |
| <b>View group</b>    | Group to which this view belongs (if any).                                                                                                                            |

## Simple tabular views

Simple tabular views offer the possibility to define criteria to filter records and also to select the columns that will be displayed in the table.

|                                     |                                                                                                                       |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>Displayed columns</b>            | Specifies the columns that will be displayed in the table.                                                            |
| <b>Sorted columns</b>               | Specifies the sort order of records in the table. See <a href="#">Sorting data</a> [p 111].                           |
| <b>Filter</b>                       | Defines filters for the records to be displayed in the table.<br>See <a href="#">Criteria editor</a> [p 291].         |
| <b>Pagination limit</b>             | Forces a limit to the number of visible records.                                                                      |
| <b>Grid edit</b>                    | If enabled, users of this view can switch to grid edit, so that they can edit records directly from the tabular view. |
| <b>Disable create and duplicate</b> | If yes, users of this view cannot create nor duplicate records from the grid edit.                                    |

## Hierarchical views

A hierarchy is a tree-based representation of data that allows emphasizing relationships between tables. It can be structured on several relationship levels called dimension levels. Furthermore, filter criteria can be applied in order to define which records will be displayed in the view.

### Hierarchy dimension

A dimension defines dependencies in a hierarchy. For example, a dimension can be specified to display products by category. You can include multiple dimension levels in a single view.

### Hierarchical view configuration options

This form allows configuring the hierarchical view options.

|                                        |                                                                                                                                                                     |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Display records in a new window</b> | If 'Yes', a new window will be opened with the record. Otherwise, it will be displayed in a new page of the same window.                                            |
| <b>Prune hierarchy</b>                 | If 'Yes', hierarchy nodes that have no children and do not belong to the target table will not be displayed.                                                        |
| <b>Display orphans</b>                 | If 'Yes', hierarchy nodes without a parent will be displayed.                                                                                                       |
| <b>Display root node</b>               | If 'No', the root node of the hierarchy will not be displayed in the view.                                                                                          |
| <b>Root node label</b>                 | Localized label of the hierarchy root node.                                                                                                                         |
| <b>Toolbar on top of hierarchy</b>     | Allows to set the toolbar on top of the hierarchy.                                                                                                                  |
| <b>Display non-matching children</b>   | In a recursive case, when a search filter is applied, allows the display of non-matching children of a matching node during a search.                               |
| <b>Remove recursive root leaves</b>    | In a recursive case, when a search filter is applied or if the mode is 'pruned', removes from the display the root leaves.                                          |
| <b>Detect cycle</b>                    | Allow cycle detection and display in a recursive case, the oldest node record will be chosen as the cycle root. Limitation: does not work in search or pruned mode. |

### Labels

For each dimension level that references another table, it is possible to define localized labels for the corresponding nodes in the hierarchy. The fields from which to derive labels can be selected using the built-in wizard.

### Filter

The criteria editor allows creating a record filter for the view.

See also [Criteria editor](#) [p 291]

### Ordering field

In order to enable specifying the position of nodes in a hierarchical view, you must designate an eligible ordering field defined in the table on which the hierarchical view is applied. An ordering field must have the 'Integer' data type and have a 'Hidden' default view mode in its advanced properties in the data model definition.

Except when the ordering field is in 'read-only' mode or when the hierarchy is filtered, any field can be repositioned.

By default, if no ordering field is specified, child nodes are sorted alphabetically by label.

#### Attention

Do not designate a field that is meant to contain data as an ordering node, as the data will be overwritten by the hierarchical view.

### Actions on hierarchy nodes

Each node in a hierarchical view has a menu ▼ containing contextual actions.

Leaf nodes can be dissociated from their parent record using 'Detach from parent'. The record then becomes an orphan node in the tree, organized under a container "unset" node.

Leaf nodes can also change parent nodes, using 'Attach to another parent'. If, according to the data model, a node can have relationships to multiple parents, the node will be both under the current parent and added under the other parent node. Otherwise, the leaf node will be moved under the other parent node.

### View sharing

Users having the 'Share views' permission on a view are able to define which users can display this view from their 'View' menu.

To do so, simply add profiles to the 'Share with' field of the view's configuration screen.

### View publication

Users having the 'Publish views' permission can publish views present in their 'View' menu.

A published view is then available to all users via Web components, workflow user tasks, data services and perspectives. To publish a view, go to *View > Manage views > name of the view > Publish*.

## 20.5 Views management

### Manage recommended views

When a user logs in with no view specified, their recommended view (if any) is applied. Otherwise, the default view is applied. The 'Manage recommended views' action allows defining assignment rules of recommended views depending on users and roles.

Available actions on recommended views are: change order of assignment rules, add a rule, edit existing rule, delete existing rule.

Thus, for a given user, the recommended views are evaluated according to the user's profile: the applied rule will be the first that matches the user's profile.

**Note**

The 'Manage recommended view' feature is only available to dataset owners.

## Manage views

The 'Manage views' sub-menu offers the following actions:

|                                               |                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Define this view as my favorite</b>        | Only available when the currently displayed view is NOT the recommended view. The favorite view will be automatically applied when accessing the table.                                                                                                                                                 |
| <b>Define recommended view as my favorite</b> | Only available when a favorite view has already been defined. This will remove the user's current favorite view. A recommended view, similarly to a favorite view, will be automatically applied when accessing the table. <b>This menu item is not displayed if no favorite view has been defined.</b> |

## 20.6 Grid edit

The grid edit feature allows to modify data in a table view. This feature can be accessed by clicking on the  button.

Accessing the grid edit from a table view requires that the feature be previously activated in the view configuration.

See also [Grid edit](#) [p 114]

### Copy/paste

The copy/paste of one or more cells into another one in the same table can be done through the *Edit* menu. It is also possible to use the associated keyboard shortcuts *Ctrl+C* and *Ctrl+V*.

This system does not use the operating system clipboard, but an internal mechanism. As a consequence, copying and pasting a cell in an external file will not work. Conversely, pasting a value into a table cell won't work either.

All simple type fields using built-in widgets are supported.

## 20.7 History

The history feature allows tracking changes on master data.

The history feature must have been previously enabled at the data model level. See [Advanced properties for tables](#) [p 59] for more information.

To view the history of a dataset, select *Actions > History* in the navigation pane.

To view the history of a table or of a selection of records, select *Actions* > *View history* in the workspace.

Several history modes exist, which allow viewing the history according to different perspectives:

|                                                         |                                                                                                  |
|---------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| <b>History in current dataspace</b>                     | The table history view displays operations on the current branch. This is the default mode.      |
| <b>History in current dataspace and ancestors</b>       | The table history view displays operations on the current branch and on all its ancestors.       |
| <b>History in current dataspace and merged children</b> | The table history view displays operations on the current branch and on all its merged children. |
| <b>History in all dataspace</b>                         | The table history view displays operations on the whole branch hierarchy.                        |

In the history view, use the **VIEW** menu in order to switch to another history mode.

See also [History](#) [p 249]

## CHAPTER 21

# Editing data

This chapter contains the following topics:

1. [Working with records in the user interface](#)
2. [Importing and exporting data](#)
3. [Restore from history](#)

## 21.1 Working with records in the user interface

Record editing takes place in the workspace portion of the user interface.

**Note**

This action is available only to authorized users in the 'Advanced perspective' or from a specifically configured perspective.

### *Creating a record*

In a tabular view, click the **+** button located above the table.

In a hierarchical view, select 'Create a record' from the menu of the parent node under which to create the new record.

Next, enter the field values in the new record creation form. Mandatory fields are indicated by asterisks.

### *Updating an existing record*

Double-click the record to update, then edit the field values in the record form.

To discard any changes in progress and restore the fields to their values before editing, click the **Revert** button.

### *Duplicating a record*

To duplicate a selected record, select **Actions > Duplicate**.

A new record creation form pre-populates the field values from the record being duplicated. The primary key must be given a unique value, unless it is automatically generated (as is the case for auto-incremented fields).

## Deleting records

To delete one or more selected records, select **Actions > Delete**.

## Comparing two records

To compare two selected records, select **Actions > Compare**.

### Note

The content of complex terminal nodes, such as aggregated lists and user defined attributes, are excluded from the comparison process. That is, the compare service ignores any differences between the values of the complex terminal nodes in the records.

## 21.2 Importing and exporting data

In a table, records can be exported to or imported from CSV or XML format.

You can either manually select certain records of the table to export, or you can export the whole table.

### See also

[CSV Services](#) [p 219]

[XML Services](#) [p 213]

## 21.3 Restore from history

When history is enabled on a table, it is possible to restore a record to a previous state, based on its registered history. If the record (identified by its primary key) still exists in the table, it will be updated with the historized values to be restored. Otherwise, it will be created.

In order to restore a record to a previous state, select a record in the history table view, and the select **Actions > Restore from history** in the workspace. A summary screen is displayed with the details of the update or creation to be performed.

The restore feature is available only on one record at a time.

If a table trigger must have a specific behavior on restore, different from the one on regular create and update, the developer can use the method `TableTriggerExecutionContext.isHistoryRestoreAPI`.

### Note

This feature has limitations linked to the limitations of the history feature:

- the 'restore from history' feature is not available on tables containing lists that are not supported by history. See [Data model limitations](#) [p 254].
- computed values, encrypted values and fields on which history has been disabled are ignored when restoring a record from history, since these fields are not historized.

See also [History](#) [p 249]



## CHAPTER 22

# Working with existing datasets

This chapter contains the following topics:

1. [Validating a dataset](#)
2. [Duplicating a dataset](#)
3. [Deactivating a dataset](#)
4. [Managing dataset permissions](#)

## 22.1 Validating a dataset

To validate a dataset at any time, select **Actions > Validate** from the navigation pane. A generated report provides the results of the validation. From the validation report, you have the option to update the reported validation status by clicking the **Revalidate** button, or to click the **Reinitialize validation report** button to clear the current validation report associated with the dataset in order to be able to rerun a full validation from scratch.

Validations of data can also be run at the table level by navigating to the desired table from the navigation pane, then selecting **Actions > Validate** from the workspace.

See [Validation](#) [p 298] for detailed information about incremental data validation.

## 22.2 Duplicating a dataset

To duplicate an existing dataset, select it from the '[Select dataset](#) [p 107]' ▾ menu in the navigation pane, then select **Actions > Duplicate**.

## 22.3 Deactivating a dataset

When a dataset is activated, it will be subject to validation. That is, all mandatory elements must be defined in order for the dataset to be valid. If a dataset is active and validated, it can be safely exported to external systems or to be used by other Java applications.

If a dataset is missing mandatory elements, it can be deactivated by setting the property 'Activated' to 'No' from **Actions > Information**.

## 22.4 Managing dataset permissions

Dataset permissions can be accessed by selecting **Actions > Permissions** in the navigation pane.

Permissions are defined using *profile* records. To define a new permissions profile, create a new record in the 'Access rights by profile' table.

See also [Profile](#) [p 21]

|                                        |                                                                                                                                                                                                                                                                                                       |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Profile</b>                         | Defines the profile to which these permissions apply.                                                                                                                                                                                                                                                 |
| <b>Restriction policy</b>              | <p>If 'Yes', indicates that when the permissions defined here are more strict than otherwise defined, these permissions are respected. This is contrary to the default where the most permissive rights defined take precedence.</p> <p>See <a href="#">Resolving user-defined rules</a> [p 281].</p> |
| <b>Dataset actions</b>                 | Specifies the permissions for actions on the dataset.                                                                                                                                                                                                                                                 |
| <b>Create a child dataset</b>          | Indicates whether the profile can create a child dataset. Inheritance also must be activated in the data model.                                                                                                                                                                                       |
| <b>Duplicate the dataset</b>           | Indicates whether the profile can duplicate the dataset.                                                                                                                                                                                                                                              |
| <b>Delete the dataset</b>              | Indicates whether the profile can delete the dataset.                                                                                                                                                                                                                                                 |
| <b>Activate/deactivate the dataset</b> | Indicates whether the profile can modify the <i>Activated</i> property in the dataset information. See <a href="#">Deactivating a dataset</a> [p 121].                                                                                                                                                |
| <b>Create a view</b>                   | Indicates whether the profile can create views and hierarchies in the dataset.                                                                                                                                                                                                                        |
| <b>Tables policy</b>                   | Specifies the default permissions for all tables. Specific permissions can also be defined for a table by clicking the '+' button.                                                                                                                                                                    |
| <b>Create a new record</b>             | Indicates whether the profile can create records in the table.                                                                                                                                                                                                                                        |
| <b>Overwrite inherited record</b>      | Indicates whether the profile can override inherited records in the table. This permission is useful when using dataset inheritance.                                                                                                                                                                  |
| <b>Occult inherited record</b>         | Indicates whether the profile can occult inherited records in the table. This permission is useful when using dataset inheritance.                                                                                                                                                                    |
| <b>Delete a record</b>                 | Indicates whether the profile can delete records in the table.                                                                                                                                                                                                                                        |
| <b>Values access policy</b>            | Specifies the default access permissions for all the nodes of the dataset and allows the definition of permissions for                                                                                                                                                                                |

specific nodes. The default access permissions are used if no custom permissions have been defined for a node.

The specific policy selector allows granting specific access permissions for a node. The links "ReadOnly", "ReadWrite", and "Hidden" set the corresponding access levels for the selected nodes.

It is possible to remove custom access permissions using the "(default)" link.

---

|                           |                                                                                                                            |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Rights on services</b> | This section specifies the access permissions for services. A service is not accessible to a profile if it is crossed-out. |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------|

---

## CHAPTER 23

# Dataset inheritance

Using the concept of dataset inheritance, it is possible to create child datasets that branch from a parent dataset. Child datasets inherit values and properties by default from the parent dataset, which they can then override if necessary. Multiple levels of inheritance can exist.

An example of using dataset inheritance is to define global default values in a parent dataset, and create child datasets for specific geographical areas, where those default values can be overridden.

**Note**

By default, dataset inheritance is disabled. It must be explicitly activated in the underlying data model.

See also [Data model configuration](#) [p 40]

This chapter contains the following topics:

1. [Dataset inheritance structure](#)
2. [Value inheritance](#)

## 23.1 Dataset inheritance structure

Once the root dataset has been created, create a child dataset from it using the **+** button in the dataset selector in the navigation pane.

**Note**

- A dataset cannot be deleted if it has child datasets. The child datasets must be deleted first.
- If a child dataset is duplicated, the newly created dataset will be inserted into the existing dataset tree as a sibling of the duplicated dataset.

## 23.2 Value inheritance

When a child dataset is created, it inherits all its field values from the parent dataset. A record can either keep the default inherited value or override it.

In tabular views, inherited values are marked in the top left corner of the cell.

The  button can be used to override a value.

## Record inheritance

A table in a child dataset inherits the records from the tables of its ancestor datasets. The table in the child dataset can add, modify, or delete records. Several states are defined to differentiate between types of records.

|                    |                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Root</b>        | A root record is a record that was created in the current dataset and does not exist in the parent dataset. A root record is inherited by the child datasets of the current dataset.                                                             |
| <b>Inherited</b>   | An inherited record is one that is defined in an ancestor dataset of the current dataset.                                                                                                                                                        |
| <b>Overwritten</b> | An overwritten record is an inherited record whose values have been modified in the current dataset. The overwritten values are inherited by the child datasets of the current dataset.                                                          |
| <b>Occulted</b>    | An occulted record is an inherited record which has been deleted in the current dataset. It will still appear in the current dataset as a record that is crossed out, but it will not be inherited in the child datasets of the current dataset. |

When the inheritance button  is toggled on, it indicates that the record or value is inherited from the parent dataset. This button can be toggled off to override the record or value. For an occulted record, toggle the button on to revert it to inheriting.

The following table summarizes the behavior of records when creating, modifying or deleting a record, depending on its initial state.

| State              | Create                                                                                                                                                                                          | Modify value                                                                                                                                                                                    | Delete                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Root</b>        | Standard new record creation. The newly created record will be inherited in child datasets of the current dataset.                                                                              | Standard modification of an existing record. The modified values will be inherited in the child datasets of the current dataset.                                                                | Standard record deletion. The record will no longer appear in the current dataset and the child datasets of the current dataset. |
| <b>Inherited</b>   | If a record is created using the same primary key as an existing inherited record, that record will be overwritten and its value will be the one submitted at creation.                         | An inherited record must first be marked as overwritten in order to modify its values.                                                                                                          | Deleting an inherited record changes its state to occulted.                                                                      |
| <b>Overwritten</b> | Not applicable. Cannot create a new record if the primary key is already used in the current dataset.                                                                                           | An overridden record can be returned to the inherited state, but its modified value will be lost.<br><br>Individual values in an overridden record can be set to inheriting or can be modified. | Deleting an overwritten record changes its state to occulted.                                                                    |
| <b>Occulted</b>    | If a record is created using the primary key of an existing occulted record, the record state will be changed to overwritten and its value modified according to the one submitted at creation. | Not applicable. An occulted record cannot be modified.                                                                                                                                          | Not applicable. An occulted record is already considered to be deleted.                                                          |

See also [Record lookup mechanism](#) [p 270]





---

# Workflow models

---

## CHAPTER 24

# Introduction to workflow models

This chapter contains the following topics:

1. [Overview](#)
2. [Using the Workflow Models area user interface](#)
3. [Generic message templates](#)
4. [Limitations of workflows](#)

## 24.1 Overview

### ***What is a workflow model?***

Workflows in TIBCO EBX® facilitate the collaborative management of data in the repository. A workflow can include human actions on data and automated tasks alike, while supporting notifications on certain events.

The first step of realizing a workflow is to create a *workflow model* that defines the progression of steps, responsibilities of users, as well as other behavior related to the workflow.

Once a workflow model has been defined, it can be validated and published as a *workflow publication*. Data workflows can then be launched from the workflow publication to execute the steps defined in the workflow model.

See also

[Workflow model \(glossary\)](#) [p 27]

[Data workflow \(glossary\)](#) [p 29]


### ***Basic concepts related to workflow models***

A basic understanding of the following terms is necessary to proceed with the creation of workflow models:

- [script task](#) [p 28]
- [user task](#) [p 28]
- [work item](#) [p 29]
- [workflow condition](#) [p 28]
- [sub-workflow invocation](#) [p 28]

- [wait task](#) [p 28]
- [data context](#) [p 28]

## 24.2 Using the Workflow Models area user interface



### Note

This area is available only to authorized users in the 'Advanced perspective'. Only authorized users can access these interfaces.

## 24.3 Generic message templates

Notification emails can be sent to inform users of specific events during the execution of data workflows.

Generic templates can be defined and reused by any workflow model in the repository. To work with generic templates, select 'Message templates' from the Workflow Models area **Actions** menu.

These templates, which are shared by all workflow models, are included statically at workflow model publication. Thus, in order to take template changes into account, you must update your existing publication by re-publishing the affected models.

Please note that, if you want to export those templates in an archive, you will have to select the dataset "configuration" as it is the one containing the message templates.

When creating a new template, two fields are required:

- **Label & Description:** Specifies the localized labels and descriptions associated with the template.
- **Message:** Specifies the localized subjects and bodies of the message.

The message template can include data context variables, such as `${variable.name}`, which are replaced when notifications are sent. System variables that can be used include:

|                                      |                                                                                                                                                                                                                                                                |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>system.time</b>                   | System time of the repository.                                                                                                                                                                                                                                 |
| <b>system.date</b>                   | System date of the repository.                                                                                                                                                                                                                                 |
| <b>workflow.lastComment</b>          | Last comment on the previous user task. (Note: this variable refers to the last user task, not the current one. Also the current task is the one on which the workflow is positioned, and it also includes the completion notification of a user task).        |
| <b>workflow.lastDecision</b>         | Last decisions made on the previous user task. (Note: this variable refers to the last user task, not the current one. Also the current task is the one on which the workflow is positioned, and it also includes the completion notification of a user task). |
| <b>user.fullName</b>                 | Full name of the notified user.                                                                                                                                                                                                                                |
| <b>user.login</b>                    | Login of the notified user.                                                                                                                                                                                                                                    |
| <b>workflow.process.label</b>        | Label of the current workflow.                                                                                                                                                                                                                                 |
| <b>workflow.process.description</b>  | Description of the current workflow.                                                                                                                                                                                                                           |
| <b>workflow.workItem.label</b>       | Label of the current work item.                                                                                                                                                                                                                                |
| <b>workflow.workItem.description</b> | Description of the current work item.                                                                                                                                                                                                                          |
| <b>workflow.workItem.offeredTo</b>   | Role to which the current work item has been offered.                                                                                                                                                                                                          |
| <b>workflow.workItem.allocatedTo</b> | User to whom the current work item has been allocated.                                                                                                                                                                                                         |
| <b>workflow.workItem.link</b>        | <p>Link to access the current work item in the work item inbox, using the Web Component API.</p> <p>This link can only be computed if a current work item is defined and if the URL is configured in Workflow-executions, in the email configuration.</p>      |

|                                                |                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>workflow.workItem.link.allocateAndStart</b> | Link to access the current work item in the work item inbox, using the Web Component API. If the target work item has not yet been started, it will be automatically allocated to and started by the user clicking the link.<br><br>This link can only be computed if a current work item is defined and if the URL is configured in Workflow-executions, in the email configuration. |
| <b>workflow.currentStep.label</b>              | Label of the current step.                                                                                                                                                                                                                                                                                                                                                            |
| <b>workflow.currentStep.description</b>        | Description of the current step.                                                                                                                                                                                                                                                                                                                                                      |

### Example

Generic template message:

*Today at \${system.time}, a new work item was offered to you*

Resulting email:

*Today at 15:19, a new work item was offered to you*

## 24.4 Limitations of workflows

The following functionality is currently unsupported in EBX®:

- **Scheduled tasks**, task executed as soon as its turn comes, and whose execution cannot be delayed.
- **Event tasks**, allowing the workflow to move forward upon receiving an event, such as a web service call.
- **Time limitation** on a task duration.

Related concepts [Data workflows](#) [p 158]

## CHAPTER 25

# Creating and implementing a workflow model

This chapter contains the following topics:

1. [Creating a workflow model](#)
2. [Implementing the steps](#)
3. [User tasks](#)
4. [Script tasks](#)
5. [Conditions](#)
6. [Sub-workflow invocations](#)
7. [Wait tasks](#)
8. [Visualizing the workflow diagram](#)

## 25.1 Creating a workflow model

A new workflow model can be created in the **Workflow Models** area. The only required information at creation is a name that is unique in the repository.

The steps of the workflow model are initialized with a single transition. In order to fully implement the workflow model, you must define the sequence of steps beyond this initial transition.

## 25.2 Implementing the steps

A workflow model defines steps that correspond to different operations that must be performed on data, and associated conditions. The following types of steps exist:

- User task
- Script task
- Condition
- Sub-workflow invocation
- Wait task

A data context is linked to each data workflow. This data context can be used to define variables that can be used as input and output variables in the steps of the workflow.

### ***Progress strategy of the next step***

For each step type (excluding sub-workflows invocations), a property is available to define which progress strategy has to be applied for the next step. Upon step completion, this strategy is evaluated in order to define the navigation when the workflow is executed. By default, the progress strategy is set to 'Display the work items table'. In that case, after the step has been executed, the work items table (work items inbox or monitoring > work items) is automatically displayed, in order to select the following work item.

Another strategy can be selected: 'Automatically open the next step'. This strategy allows the user to keep working on this workflow and to directly execute the next step. If, following to this execution, a work item is reached and the connected user can start it, then the work item is automatically opened (if several work items are reached, the first created is opened). Otherwise, the next step progress strategy is evaluated. If no work item has been reached, the work items table will be displayed.

This strategy is used to execute several steps in a row without going back to the work items inbox.

There are some limitations that will lead to disregard this strategy. In that case, the work items table is automatically displayed. This property will be disregarded when: the next step is a sub-workflow; or the current step is a user task with more than one work item.

In the case of conditions, two other strategies are available: 'If true, automatically open the next step' and 'If false, automatically open the next step'. These strategies allow choosing which strategy will be applied according to the condition result.

### ***Hidden in graphical view***

For each step type, a property is available to define which steps must be hidden in the workflow progress view by default.

If this property is enabled, the step will be automatically hidden in the workflow progress view for non-administrator users (neither built-in administrator nor workflow administrator). Hidden steps can be displayed on demand.

## **25.3 User tasks**

User tasks are steps that involve actions to be performed by human users. Their labels and descriptions can be localized.

### ***Mode***

For backward compatibility reasons, two user task modes are available: the default mode and the legacy mode.

By default, a user task generates a single work item. This mode offers more features, such as offering a work item to a list of profiles or directly displaying the avatars in the workflow progress view.

In the legacy mode, a user task can generate several work items.

By default, the user task creation service is hidden in legacy mode. To display it, a property should be enabled in the `ebx.properties` file. For more information, see [Disabling user task legacy mode](#) [p 363].

### ***List of profiles***

The definition of the profiles of a user task may vary depending on the user task mode.



**[Default] Offered to the following profiles**

The defined profiles are the roles or the users to whom the user task is being offered. When executing the user task, a single work item is generated. If a single user is defined, the work item is automatically assigned to this user. If a role is defined, the work item is offered to the members of the role. If several users and roles are defined, the work item is offered simultaneously to these users and to the members of these roles.

**[Legacy mode] Participants**

The participants are the roles or the users to whom the user task is intended. By default, when executing the user task, a work item is generated by profile. If a profile refers to a user instead of a role, the work item is directly allocated to that user. If a profile is a role, the work item is offered to the members of the role.

**For more information**

See also [Extension](#) [p 139]

**Service**

TIBCO EBX® includes the following built-in services:

- Access a dataspace
- Access data (default service)
- Access the dataspace merge view
- Compare contents
- Create a new record
- Duplicate a record.
- Export data to a CSV file
- Export data to an XML file
- Import data from a CSV file
- Import data from an XML file
- Merge a dataspace
- Validate a dataspace, a snapshot or a dataset

See also [EBX® built-in services](#) [p 199]

**Configuration****Main options > Enable reject**

By default, only the *accept* action is offered to the user when saving a decision.

It is possible to also allow users to reject the work item by setting this field to 'Yes'.

## Main options > Enable confirmation request

By default, a confirmation request is displayed after user task execution when the user saves the decision by clicking the 'Accept' or 'Reject' button.

To disable this confirmation prompt, set this field to 'Yes'.

## Main options > Enable comments

By default, comments are enabled. When a work item is open, a 'Comments' button is displayed and allows the user to enter a comment.

It is possible to hide this 'Comments' button by setting this property to *No*.

## Main options > Comments required

By default, it is optional to submit a comment associated with a work item.

It is possible to require the user to enter a comment before saving the decision by setting this field to the desired validation criteria. Comments can be set to be always required, required only if the work item has been accepted, or required only if the work item has been rejected.

## Main options > Customized labels

When the user task is run, the user can accept or reject the work item by clicking the corresponding button. In the workflow model, it is possible for a user task to define a customized label and confirmation message for these two buttons. This can be used to adapt the buttons to a more specific context.

## [Legacy mode] Termination > Task termination criteria

A single user task could be assigned to multiple *participants* and thus generate multiple work items during workflow execution. When defining a user task in the workflow model, you can select one of the predefined methods for determining whether the user task is finished, based on the statuses of its component work items. When the user task's exit requirement has been satisfied, the data workflow will move on to the next step defined in its model.

For example, for the case of a user task where a product record needs to be approved, you could designate three potential participants. The task termination criteria can specify whether the product record needs to be approved by all three users, or only the first user to respond.

The default user task termination criteria is 'When all work items have been accepted.'

### Note

If you specify a service extension overriding the method `UserTask.handleWorkItemCompletion` to handle the determination of the user task's completion, it is up to the developer of the extension to verify from within their code that the task termination criteria defined through the user interface has been met. See `UserTask.handleWorkItemCompletionAPT` in the JavaDoc for more information

## [Legacy mode] Termination > Reject tolerance

By default, if a user rejects a work item during workflow execution, the user task is placed into an error state and the workflow progress is halted. When the user task is in an error state, a workflow administrator must intervene by replaying the step where the error occurred in order to continue the workflow execution.

In order to change this default behavior, it is possible to define a certain number of work item rejections to tolerate. While within the limit of tolerated rejections, no error will occur and it is the task termination criteria that determines when to end the user task.

The following task termination criteria automatically tolerate all rejections:

- 'When all work items have been either accepted or rejected'
- 'Either when all work items have been accepted, or as soon as one work item has been rejected'

## Extension

A custom class can be specified in order for the task to behave dynamically in the context of a given data workflow. For example, this can be used to create work items or complete user tasks differently than the default behavior.

The specified rule is a JavaBean that must extend the `UserTaskAPI` class.

### Attention

If a rule is specified and the `handleWorkItemCompletion` method is overridden, the completion strategy is no longer automatically checked. The developer must check for completion within this method.

## Notification

A notification email can be sent to users when specific events occur. For each event, you can specify a content template.

It is possible to define a monitor profile that will receive all emails that are sent in relation to the user task.

See also [Generic message templates](#) [p 131]

## Reminder

Reminder emails for outstanding offered or allocated work items can be periodically sent to the concerned users. The recipients of the reminder are the users to whom the work item is offered or allocated, as well as the recipients on copy.

The content of the reminder emails is determined by the current state of the work item. That is, if the work item is offered, the notification will use the "Offered work items" template; if the work item is allocated, the notification will use the "Allocated work items" template.

## Deadline

Each user task can have a completion deadline. If this date passes and associated works items are not completed, a notification email is sent to the concerned users. This same notification email will then be sent daily until the task is completed.

There are two deadline types:

- *Absolute deadline*: A calendar date.
- *Relative deadline*: A duration in hours, days or months. The duration is evaluated based on the reference date, which is the beginning of the user task or the beginning of the workflow.

## 25.4 Script tasks

Script tasks are automatic tasks that are performed without human user involvement.

Two types of script tasks exist, which, once defined, can be used in workflow model steps:

|                             |                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Library script task</b>  | <p>EBX® includes a number of built-in library script tasks, which can be used as-is.</p> <p>Any additional library script tasks must be declared in a <code>module.xml</code> file. A library script task must define its label, description and parameters. When a user selects a library script task for a step in a workflow model, its associated parameters are displayed dynamically.</p> |
| <b>Specific script task</b> | <p>Specifies a Java class that performs custom actions. The associated class must belong to the same module as the workflow model. Its labels and descriptions are not displayed dynamically to users in workflow models.</p>                                                                                                                                                                   |

[Packaging TIBCO EBX® modules](#) [p 461]

### *Library script tasks*

EBX® includes the following built-in library script tasks:

- Create a dataspace
- Create a snapshot
- Merge a dataspace
- Import an archive
- Close a dataspace
- Set a data context variable
- Send an email
- Delete records (Note: this script can remove several records)

Library script tasks are classes that extend the class `ScriptTaskBeanAPI`. Besides the built-in library script tasks, additional library script tasks can be defined for use in workflow models. Their labels and descriptions can be localized.

The method `ScriptTaskBean.executeScriptAPI` is called when the data workflow reaches the corresponding step.

#### **Attention**

The method `ScriptTaskBean.executeScriptAPI` must not create any threads because the data workflow moves on as soon as the method is executed. Each operation in this method must therefore be synchronous.

See the [example](#) [p 552].

It is possible to dynamically set variables of the library script task if its implementation follows the Java Bean specification. Variables must be declared as parameters of the bean of the library script task in `module.xml`. The workflow data context is not accessible from a Java bean.

#### Note

Some built-in library script tasks are marked as "deprecated" because they are not compatible with internationalization. It is recommended to use the new script tasks that are compatible with internationalization.

## Specific script tasks

Specific script tasks are classes that extend the class [Sample of ScriptTask](#) [p 552].

The method `ScriptTask.executeScriptAPI` is called when the data workflow reaches the corresponding step.

#### Attention

The method `ScriptTask.executeScriptAPI` must not create any threads because the data workflow moves on as soon as the method is executed. Each operation in this method must therefore be synchronous.

See the [example](#) [p 552].

It is not possible to dynamically set the variables of the bean for specific script tasks. However, the workflow data context is accessible from the Java bean.

## 25.5 Conditions

Conditions are decision steps in workflows.

Two types of conditions exist, which, once defined, can be used in workflow model steps:

|                           |                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Library condition</b>  | <p>EBX® includes a number of built-in library conditions, which can be used as-is.</p> <p>Any additional library script tasks must be declared in a <code>module.xml</code> file. A library condition must define its label, description and parameters. When a user selects a library condition for a step in a workflow model, its associated parameters are displayed dynamically.</p> |
| <b>Specific condition</b> | <p>Specifies a Java class that implements a custom condition. The associated class must belong to the same module as the workflow model. Its labels and descriptions are not displayed dynamically to users in workflow models.</p>                                                                                                                                                       |

[Packaging TIBCO EBX® modules](#) [p 461]

### Library conditions

EBX® includes the following built-in library conditions:

- Dataspace modified?
- Data valid?
- Last user task accepted?
- Value null or empty?
- Values equals?

Library conditions are classes that extend the class `ConditionBeanAPI`. Besides the built-in library conditions, additional library conditions can be defined for use in workflow models. Their labels and descriptions can be localized.

See the [example](#) [p 555].

It is possible to dynamically set variables of the library condition if its implementation follows the Java Bean specification. Variables must be declared as parameters of the bean of the library condition in `module.xml`. The workflow data context is not accessible from a Java bean.

### ***Specific conditions***

Specific conditions are classes that extend the class `ConditionAPI`.

See the [example](#) [p 554].

It is not possible to dynamically set the variables of the bean for specific conditions. However, the workflow data context is accessible from the Java bean.

## **25.6 Sub-workflow invocations**

Sub-workflow invocation steps in workflow models put the current workflow into a waiting state and invoke one or more workflows.

It is possible to include another workflow model definition in the current workflow by invoking it alone in a sub-workflow invocation step.

If multiple sub-workflows are invoked by a single step, they are run concurrently, in parallel. All sub-workflows must be terminated before the original workflow continues onto the next step. The label and description of each sub-workflow can be localized.

Two types of sub-workflow invocations exist:

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Static</b>  | <p>Defines one or more sub-workflows to be invoked each time the step is reached in a data workflow. For each sub-workflow, it is possible to set its localized labels and descriptions, as well as the input and output variable mappings in its data context.</p> <p>This mode is useful when the sub-workflows to be launched and the output mappings are predetermined.</p>                                                                                                                                                                                                      |
| <b>Dynamic</b> | <p>Specifies a Java class that implements a custom sub-workflow invocation. All workflows that could be potentially invoked as sub-workflows by the code must be declared as dependencies.</p> <p>The workflow data context is directly accessible from the Java bean.</p> <p>Dynamic sub-workflow invocations must be declared in a <code>module.xml</code> file.</p> <p>This mode is useful when the launch of sub-workflows is conditional (for example, if it depends on a data context variable), or when the output mapping depends on the execution of the sub-workflows.</p> |

## 25.7 Wait tasks

Wait task steps in workflow models put the current workflow into a waiting state until a specific event is received.

When a wait task is reached, the workflow engine generates a unique resume identifier associated with the wait task. This identifier will be required to resume the wait task, and as a consequence the associated workflow.

A wait task specifies which profile is authorized to resume the wait task; and a Java class that implements a wait task bean: `WaitTaskBeanAPT`.

The workflow data context is directly accessible from the Java bean.

Wait task beans must be declared in a `module.xml` file.

First, the wait task bean is called when the workflow starts waiting. At this time, the generated resume identifier is available to call a web service for example. Then, the wait task bean is called when the wait task is resumed. In this way, the data context may be updated according to the received parameters.

### Note

The built-in administrator always has the right to resume a workflow.

## 25.8 Visualizing the workflow diagram

### About

Once a workflow model is defined, one can view the model in a BPMN-like diagram.

This service is available with a dedicated button on the toolbar of the hierarchical view. The icon is the following: .

This service provides a view with limited edition capabilities: it is only possible to modify existing steps, but links edition and creation of steps still need to be done through the hierarchical view. This diagram can help modelers have a clear view of the workflow model they are designing.

Please also note that, although the diagram is derived from BPMN standards, it is not a strict representation of BPMN since EBX® workflow concepts are slightly different.

### Saving the layout

It is possible to save the modified layout. Please note that this is not a user-based save: it will be shared by all the users.

### Actions

|                      |                        |
|----------------------|------------------------|
| <b>Export as PNG</b> | Creates a PNG image.   |
| <b>Export as SVG</b> | Creates an SVG image.  |
| <b>Export as PDF</b> | Creates a one-page PDF |

### View

|                                   |                                                                |
|-----------------------------------|----------------------------------------------------------------|
| <b>Layout &gt; Default layout</b> | Applies the default layout to the diagram.                     |
| <b>Grid &gt; Show/Hide grid</b>   | Shows the grid if the grid is not visible, hides it otherwise. |

### Buttons

|                              |                                                        |
|------------------------------|--------------------------------------------------------|
| <b>Save layout</b>           | Saves the current layout.                              |
| <b>Save layout and close</b> | Saves the current layout and closes the service.       |
| <b>Revert</b>                | Reverts changes and reloads a previously saved layout. |
| <b>Close</b>                 | Closes the service.                                    |



## ***Additional features***

The diagram view offers useful additional features

|                                  |                                                                                                                                                                                                                    |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Undo last action</b>          | CTRL + Z                                                                                                                                                                                                           |
| <b>Zoom in/Zoom out.</b>         | Mouse middle button then mouse wheel / CTRL then mouse wheel.                                                                                                                                                      |
| <b>Multiple selection</b>        | Click on the nodes or links selected holding down the CTRL button / Draw a selection rectangle (you will need to hold down the left click for 1 second before drawing the area).                                   |
| <b>Customizing links drawing</b> | When clicking on a link, you can either move the segments by dragging the squares which appear on the corners, or separate a specific segment by moving the circle in the middle.                                  |
| <b>Edit a step</b>               | Double clicking on a step will display an edition form.                                                                                                                                                            |
| <b>Overview</b>                  | A panel is now available with a miniature workflow diagram view which can be used to navigate within it. This panel can be collapsed, expanded and dragged inside the area allocated to the workflow diagram view. |



## CHAPTER 26

# Configuring the workflow model

This chapter contains the following topics:

1. [Information associated with a workflow model](#)
2. [Workflow model properties](#)
3. [Data context](#)
4. [Custom workflow execution views](#)
5. [Permissions on associated data workflows](#)
6. [Workflow model snapshots](#)
7. [Deleting a workflow model](#)

## 26.1 Information associated with a workflow model

To view and edit the owner and documentation of your workflow model, select 'Information' from the [workflow model 'Actions'](#) [p 131] menu for your workflow model in the navigation pane.

|                                |                                                                                                                                                      |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Owner</b>                   | Specifies the workflow model owner, who will have the rights to edit the workflow model's information and define its permissions.                    |
| <b>Localized documentation</b> | Localized labels and descriptions for the workflow model.                                                                                            |
| <b>Activated</b>               | <i>This property is deprecated.</i> Whether the workflow model is activated. A workflow model must be activated in order to be able to be published. |

## 26.2 Workflow model properties

Configuration for a workflow model is accessible in the navigation pane under 'Workflow model configuration'.

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Module name</b>                       | Module containing specific Java resources (user task extensions, specific scripts and conditions).                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Notification of start</b>             | The list of profiles to which to send notifications, based on a template, when a data workflow is launched.<br>See <a href="#">Generic message templates</a> [p 131].                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Notification of completion</b>        | The list of profiles to which to send notifications, based on a template, when a data workflow is completed. The notification is only sent if the workflow has been completed under normal circumstances, that is, not due to an administration action.<br>See <a href="#">Generic message templates</a> [p 131].                                                                                                                                                                                                                                     |
| <b>Notification of error</b>             | The list of profiles that will receive notifications, based on a template, when a data workflow is in error state.<br>See <a href="#">Generic message templates</a> [p 131].                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Priority</b>                          | By default, each workflow associated with this model will be launched with this priority. Setting a priority is optional. If no priority is defined here, and a default priority is set for the repository, the repository default priority will be used for any associated workflow with no priority assigned. See <a href="#">Work item priorities</a> [p 169] for more information.<br><br><b>Note:</b> Only users who are defined as workflow administrators will be able to manually modify the priority level of any associated data workflows. |
| <b>Activate quick launch</b>             | By default, when a workflow is launched, the user is prompted to enter a documentation for the new workflow in an intermediate form. This documentation is optional. Setting the 'Activate quick launch' property to 'Yes' allows skipping this documentation step and proceeding directly to the workflow launch.                                                                                                                                                                                                                                    |
| <b>Automatically open the first step</b> | Allows determining the navigation after a workflow is launched. By default, once a workflow is launched, the current table (workflow launchers or monitoring > publications) is automatically displayed.                                                                                                                                                                                                                                                                                                                                              |

Enabling this property will allow the workflow user to keep working on the launched workflow. If, after the first workflow step is executed, a work item is reached, and this work item can be started by the workflow creator, then the work item is automatically opened (if several work items are reached, the first created is opened). This will save the user from selecting the corresponding work item from the work items inbox.

If no work item has been reached, the next step progress strategy is evaluated.

If no work item has been opened, the table from which the workflow has been launched is displayed.

**Limitation:** This property will be ignored if the first step is a sub-workflow invocation.

|                                        |                                                                                                                                                                                                    |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Workflow trigger</b>                | Component that intercepts the main events of a workflow.<br>This bean must be declared in a <code>module.xml</code> file. See the <a href="#">example</a> [p 558].                                 |
| <b>Permissions</b>                     | Permissions on actions related to the data workflows associated with the workflow model.<br>This bean must be declared in a <code>module.xml</code> file. See the <a href="#">example</a> [p 557]. |
| <b>Programmatic action permissions</b> | Defines a custom component that handles the permissions of the workflow. If set, this overrides all permissions defined in the property 'Permissions'.                                             |

## 26.3 Data context

The data context configuration can be accessed from the navigation pane.

Each workflow has its own data context, thus allowing to have its own local dataspace during its execution. This gives the possibility to store and to vary values that will direct the workflow execution.

The data context is defined by a list of variables. Each variable has the following properties:

|                         |                                                                                                                                                                                              |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>             | Identifier of the variable.                                                                                                                                                                  |
| <b>Default value</b>    | If defined, the variable will be initialized with this default value.                                                                                                                        |
| <b>Input parameter</b>  | 'Yes' must be checked in order to define this variable as an input parameter.                                                                                                                |
| <b>Output parameter</b> | 'Yes' must be checked in order to define this variable as an output parameter. Else, this variable will not be displayed in the list of output parameters, in the task definition interface. |

## 26.4 Custom workflow execution views

The workflow execution views customization can be accessed from the navigation pane.

The customization allows configuring the specific columns of the work items and workflow views (inbox, work items monitoring, active workflows monitoring and completed workflows). For each specific column, it is possible to associate an expression that can contain data context variables that will be evaluated upon display of the workflow.

## 26.5 Permissions on associated data workflows

|                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Workflow administration</b>                                 | Defines the profile that is allowed to perform administration actions on the workflows. The administration actions include the following: replay a step, resume a workflow, terminate a workflow, disable a publication and unpublish. In order to perform these actions, this profile is automatically granted the "Visualize workflows" permission. The built-in administrator always has the workflow administration rights.                                                                                                                                                            |
| <b>Workflow administration &gt; Replay a step</b>              | Defines the profile that is allowed to replay a workflow step. In order to perform this action, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Active workflows" section is available to replay a step. A profile with the "Workflow administration" permission is automatically allowed to perform this specific action. The built-in administrator always has the rights to replay a step.                                                                                                                                    |
| <b>Workflow administration &gt; Terminate workflow</b>         | Defines the profile that is allowed to terminate and clean a workflow. In order to perform this action, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Active workflows" section is available to terminate and clean an active workflow. A button in the "Completed workflows" section is available to delete a completed workflow. A profile with the "Workflow administration" permission is automatically allowed to perform this specific action. The built-in administrator always has the rights to terminate a workflow. |
| <b>Workflow administration &gt; Force a workflow to resume</b> | Defines the profile that is allowed to force resuming a waiting workflow. In order to perform this action, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Active workflows" section is available to resume a workflow. A profile with the "Workflow administration" permission is automatically allowed to perform this specific action. The built-in administrator always has the right to resume a workflow.                                                                                                                  |
| <b>Workflow administration &gt; Disable a publication</b>      | Defines the profile that is allowed to disable a workflow publication. In order to perform this action, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Publications" section is available to disable a publication. It is only                                                                                                                                                                                                                                                                                                  |

displayed on active publications. A profile with the "Workflow administration" permission is automatically allowed to perform this specific action. The built-in administrator always has the rights to disable a publication.

---

**Workflow administration >  
Unpublish**

Defines the profile that is allowed to unpublish a workflow publication. In order to perform this action, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Publications" section is available to unpublish disabled publications only. A profile with the "Workflow administration" permission is automatically allowed to perform this specific action. The built-in administrator always has the unpublish rights.

---

**Allocation management**

Defines the profile that is allowed to manage work items allocation. The allocation actions include the following: allocate work items, reallocate work items and deallocate work items. In order to perform these actions, this profile is automatically granted the "Visualize workflows" permission. The built-in administrator always has the allocation management rights.

---

**Allocation management >  
Allocate work items**

Defines the profile that is allowed to allocate work items. In order to perform these actions, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Work items" section is available to allocate a work item. It is only displayed on offered work items. A profile with the "Allocation management" permission is automatically allowed to perform this specific action. The built-in administrator always has the work items allocation rights.

---

**Allocation management >  
Reallocate work items**

Defines the profile that is allowed to reallocate work items. In order to perform this action, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Work items" section is available to reallocate a work item. It is only displayed on allocated work items. A profile with the "Allocation management" permission is automatically allowed to perform this specific action. The built-in administrator always has the work items reallocation rights.

---

**Allocation management >  
Deallocate work items**

Defines the profile that is allowed to deallocate work items. In order to perform this action, this profile is automatically granted the "Visualize workflows" permission. A button in the "Monitoring > Work items" section is available to deallocate a work item. It is only displayed on allocated work items. A profile with the "Allocation management" permission is automatically allowed to perform this specific



action. The built-in administrator always has the work items deallocation rights.

|                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Launch workflows</b>                                               | Defines the profile that is allowed to manually launch new workflows. This permission allows launching workflows from the active publications of the "Workflow launchers" section. The built-in administrator always has the launch workflows rights.                                                                                                                                                                                                                                                                      |
| <b>Visualize workflows</b>                                            | Defines the profile that is allowed to visualize workflows. By default, the end-user can only see work items that have been offered or allocated to him in the "Inbox" section. This permission also allows visualizing the publications, workflows and work items associated with this workflow model in the "Monitoring" and "Completed workflows" sections. This profile is automatically granted the "Visualize completed workflows" permission. The built-in administrator always has the visualize workflows rights. |
| <b>Visualize workflows &gt; The workflow creator can visualize it</b> | If enabled, the workflow creator has the permission to view the workflows he has launched. This restricted permission grants access to the workflows he launched and to the associated work items in the "Monitoring > Active workflows", "Monitoring > Work items" and "Completed workflows" sections. The default value is 'No'.                                                                                                                                                                                         |
| <b>Visualize workflows &gt; Visualize completed workflows</b>         | Defines the profile that is allowed to visualize completed workflows. This permission allows visualizing completed workflows in the "Completed workflows" section and accessing their history. A profile with the "Visualize workflows" permission is automatically allowed to perform this action. The built-in administrator always has the visualize completed workflows rights.                                                                                                                                        |

#### Note

A user who has no specific privileges assigned can only see work items associated with this workflow that are offered or allocated to that user.

See also [Workflow administration](#) [p 173]

## 26.6 Workflow model snapshots

The history of workflow model snapshots can be managed from **Actions > History**.

The history table displays all snapshots which contain the selected workflow model and indicates if a workflow model is published. For each snapshot, the **Actions** button allows you to export or view the corresponding workflow model.

## 26.7 Deleting a workflow model

Workflow model can be deleted, however any associated publications remain accessible in the Data Workflows area. If a new workflow model is created with the same name as a deleted workflow model, publishing will prompt to replace the old publication.

See also [Publishing workflow models](#) [p 155]

## CHAPTER 27

# Publishing workflow models

This chapter contains the following topics:

1. [About workflow publications](#)
2. [Publishing and workflow model snapshots](#)
3. [Sub-workflows in publications](#)

## 27.1 About workflow publications

Once a workflow model is defined, it must be published in order to enable authorized users to launch associated data workflows. This is done by clicking the **Publish** button in the navigation pane.

If no sub-workflow invocation steps are included in the current workflow model, you have the option of publishing other workflow models at the same time on the publication page. If the current workflow model contains sub-workflow invocation steps, it must be published alone.

Workflow models can be published several times. A publication is identified by its publication name

## 27.2 Publishing and workflow model snapshots

When publishing a workflow model, a snapshot is taken of its current state. A label and a description can be specified for the snapshot to be created. The default snapshot label is the date and time of the publication. The default description indicates the user who published the workflow model.

For each workflow model being published, the specified publication name must be unique. If a workflow model has already been published, it is possible to update an existing publication by reusing the same publication name. The names of existing workflow publications associated with a given workflow model are available in a drop-down menu. In the case of a publication update, the old version is no longer available for launching data workflows, however it will be used to terminate existing workflows. The content of different versions can be viewed in the workflow model snapshot history.

See also [Workflow model snapshots](#) [p 153]

## 27.3 Sub-workflows in publications

When publishing a workflow model containing sub-workflow invocation steps, it is not necessary to separately publish the models of the sub-workflows. From an administration standpoint, the model of the main workflow (the one currently published by a user) and the models of the sub-workflows are published as a single entity.

The system computes the dependencies to workflow models used as sub-workflows, and automatically creates one publication for each dependent model. These technical publications are dedicated to the workflow engine to launch sub-workflows, and are not available in the Workflow Data area.

The multiple publication is not available for a workflow model containing sub-workflow invocation steps. This is why the first step of the publication (selection of workflow models to publish) is not offered in this case.

Republishing the main workflow model automatically updates the invoked sub-workflow models.

Although a sub-workflow model can be published separately as a main workflow model, this will not update the version used by an already published main workflow model using this sub-workflow.

---

# Data workflows

---

## CHAPTER 28

# Introduction to data workflows

This chapter contains the following topics:

1. [Overview](#)

## 28.1 Overview

A data workflow is an executed step-by-step data management process, defined using a workflow model publication. It allows users, as well as automated procedures, to perform actions collaboratively on a set of data. Once a workflow model has been developed and published, the resulting publication can be used to launch a data workflow to execute the defined steps.

Depending on the workflow user permissions defined by the workflow model, a user may perform one or more of the following actions on associated data workflows:

- As a user with default permissions, work on and complete an assigned work item.
- As a user with workflow launching permissions, create a new data workflow from a workflow model publication.
- As a workflow monitor, follow the progress of ongoing data workflows and consult the history of completed data workflows.
- As a manager of work item allocation, modify work item allocations manually for other users and roles.
- As a workflow administrator, perform various administration actions, such as replaying steps, terminating workflows in progress, or rendering publications unavailable for launching data workflows.

### See also

[Work items](#) [p 165]

[Launching and monitoring data workflows](#) [p 171]

[Administration of data workflows](#) [p 173]

[Permissions on associated data workflows](#) [p 151]

**Related concepts** [Workflow models](#) [p 130]

## CHAPTER 29

---


# Using the Data Workflows area user interface

This chapter contains the following topics:

1. [Navigating within the interface](#)
2. [Navigation rules](#)
3. [Custom views](#)
4. [Specific columns](#)
5. [Filtering items in views](#)
6. [Graphical workflow view](#)

## 29.1 Navigating within the interface

Data workflow functionality is located in the **Data Workflows** area of the TIBCO EBX® user interface.



### Note

This area is available only to authorized users in the 'Advanced perspective' or from a specifically configured perspective. Only authorized users can access these interfaces.



The navigation pane is organized into several entries. These entries are displayed according to their associated global permission. The different entries are:

|                            |                                                                                                                                                                                                                                                                                                           |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Work items inbox</b>    | All work items either allocated or offered to you, for which you must perform the defined task.                                                                                                                                                                                                           |
| <b>Workflow launchers</b>  | List of workflow model publications from which you are allowed to launch data workflows, according to your user permissions.                                                                                                                                                                              |
| <b>Monitoring</b>          | Monitoring views on the data workflows for which you have the necessary viewing permissions.                                                                                                                                                                                                              |
| <b>Publications</b>        | Publications for which you have the necessary viewing permissions. If you have additional administrative permissions, you can also disable the ability to launch data workflows from specific publications from this view.                                                                                |
| <b>Active workflows</b>    | Data workflows in the process of execution for which you have the necessary viewing permissions. If you have additional administrative permissions, you can also perform actions such as replaying steps of data workflows, and terminating the execution of data workflows from this view.               |
| <b>Work items</b>          | Work items for which you have the necessary viewing permissions. If you have additional administrative permissions, you can also perform actions relevant to work item administration, such as allocating work items to specific users or roles from this view.                                           |
| <b>Completed workflows</b> | Data workflows that have completed their execution, for which you have the necessary viewing permissions. You can view the history of the executions of the data workflows. If you have additional administrative permissions, you can also clean completed workflows from the repository from this view. |

#### Note

Each section can be accessed through Web Components, for example, for portal integration, or programmatically using the `ServiceKey` class in the Java API.

#### See also

[Using TIBCO EBX® as a Web Component](#) [p 191]

`ServiceKey`<sup>API</sup>

## 29.2 Navigation rules

### ***Work items inbox***

By default, once a work item has been executed, the work items inbox is displayed.

This behavior can be modified according to the next step progress strategy, which can allow to execute several steps in a row without going back to the work items inbox.

See the [progress strategy of a workflow step](#) [p 136] in workflow modeling.

### ***Workflow launchers***

By default, once a workflow has been launched, the workflow launchers table is displayed.

This behavior can be modified according to the model configuration, which can allow to directly open the first step without displaying the workflow launchers table.

See [the automatic opening of the first workflow step](#) [p 148] in workflow modeling.

## 29.3 Custom views

It is possible to define views on workflow tables and to benefit from all associated mechanisms (publication included).

Permissions to create and manage workflow table views are the same as the permissions for data table views. It may thus be necessary to change the permissions in the 'Administration' section in order to benefit from this feature, by selecting *Workflow management > Workflows*.

See the [Views](#) [p 113] for more information.


## 29.4 Specific columns

By default, specific columns are hidden in the views that can benefit from it (inbox, work items monitoring, active workflows monitoring and completed workflows).

A custom view should be created and applied in order to display the specific columns. For each work item or workflow, the matching defined in the associated workflow model is then applied. If an expression is defined for a column and contains data context variables, these variables are evaluated upon display. If the expression contains built-in expressions which depend on the locale, the expression is evaluated in the default locale.

## 29.5 Filtering items in views

In certain tables, such as the 'Work item inbox', you can narrow down the entries displayed in the tables by viewing only the entries of a certain state. In such views, there is a menu from which you can select a state to see its corresponding items.



## 29.6 Graphical workflow view

Whether as a user with a work item to perform, or as a data workflow monitor or administrator, you can view the progress or the history of a data workflow execution by clicking the 'Preview'  button that appears in the 'Data workflow' column of tables throughout the data workflows user interface. This opens a pop-up displaying an interactive graphical view of the data workflow's execution. In this view, you can see the overall progress of the execution, as well as click on an individual step to view the details of its information.

If steps have been defined as hidden in the workflow modeling, they are automatically hidden in the workflow progress view for non-administrator users (non built-in administrators and non workflow administrators). A button is available to display hidden steps. The choice of users (show or hide steps) is saved by user, by publication during the user session.

For user tasks performed using the new mode (single work item), the main information about the single work item is directly displayed in the workflow progress view, when applicable: the avatar of the user associated with the work item, and the decision that has been taken for the work item (accepted or rejected).



## CHAPTER 30

---

# Work items

This chapter contains the following topics:

1. [About work items](#)
2. [Working on work items as a participant](#)
3. [Work item priorities](#)

### 30.1 About work items

A work item is a unit of work that must be performed by a human user as a part of a user task. By default, when a workflow model defines a user task, data workflows that are launched from that model's publications will generate an individual work item for each of the participants listed in the user task.

See also [Overview](#) [p 565]

#### **Work item states**

When the data workflow spawns a work item for a modelled user task during execution, the work item passes through several possible states: offered, allocated, started, and completed.

#### **Creation of work items**

##### **Default mode**

By default, a single work item is generated regardless of the list of defined profiles.

By default, if a single user is defined in the list of profiles, the created work item is in the *allocated* state.

By default, in other cases, the created work item is in the *offered* state.

##### **Note**

The default behavior described above can be overridden by a programmatic extension defined in the user task. In this case, work items may be generated programmatically and not necessarily based on the user task's list of participants.

##### **Legacy mode**

By default, for each user defined as a participant of the user task, the data workflow creates a work item in the *allocated* state.

By default, for each role defined as a participant of the user task, the data workflow creates a work item in the *offered* state.

**Note**

The default behavior described above can be overridden by a programmatic extension defined in the user task. In this case, work items may be generated programmatically and not necessarily based on the user task's list of participants.

**Variations of the work item states**


When the work item is in the *allocated* state, the defined user can directly start working on the allocated work item with the 'Take and start' action. The work item's state becomes *started*.


When the work item is in the *offered* state, any user or member of the roles to whom the work item is offered can take the work item with the 'Take and start' action'. The work item's state becomes *started*.

Before a user has claimed the offered work item, a workflow allocation manager can intervene to manually assign the work item to a single user, thus moving the work item to the *allocated* state. Then, when that user begins work on the work item by performing the action 'Start work item', the work item progresses to the *started* state.

Finally, after the user who started the work item has finished the assigned action, the concluding accept or reject action moves the work item to the *completed* state. Once a user completes a work item, the data workflow automatically progresses onto the next step defined in the workflow model.

## Diagram of the work item states





## 30.2 Working on work items as a participant

All work items relevant to you as a user (either offered or allocated to you), appear in your work items inbox. When you begin working on a work item, you can add an associated comment that will be visible to other participants of the data workflow, as well as administrators and monitors of the workflow. As long as you are still working on the work item, you can go back and edit this comment. After you have performed all the necessary actions assigned for the work item, you must signal its completion by clicking either the **Accept** or **Reject** button. The labels of these two buttons may differ depending on the context of the work item.

To review the current progress of a data workflow for which you have a waiting work item in your work item inbox, click its 'Preview'  button in the 'Data workflow' column of the table. A pop-up will show an interactive graphical view of the data workflow up until this point and the upcoming steps. You can view the details of a step by clicking on that step.

### Note

If you interrupt the current session in the middle of a started work item, for example by closing the browser or by logging out, the current work item state is preserved. When you return to the work item, it continues from the point where you left off.



## 30.3 Work item priorities

Work items may carry a priority value, which can be useful for sorting and filtering outstanding work items. The priority of a work item is set at the level of its data workflow, rather than being specific to the individual work item itself. Thus, if a data workflow is considered urgent, all its associated open work items are also considered to be urgent. By default, there are six priority levels ranging from 'Very low' to 'Urgent', however the visual representation and naming of the available priority levels depend on the configuration of your TIBCO EBX® repository.

See also [user task \(glossary\)](#) [p 28]

Related concepts [User tasks](#) [p 136]



## CHAPTER 31

# Launching and monitoring data workflows

This chapter contains the following topics:

1. [Launching data workflows](#)
2. [Monitoring activities](#)
3. [Managing work item allocation](#)

## 31.1 Launching data workflows

If a workflow model has given you the permissions to launch data workflows from its publications, you can create new data workflows from the 'Workflow launchers' entry in the navigation pane. To create a new data workflow from a workflow model publication, click the **Launch** button in the entry of the publication.

You can optionally define localized labels and descriptions for the new data workflow you are launching.

## 31.2 Monitoring activities

If a workflow model's permissions have configured your user or role for workflow monitoring, you have the ability to follow the progress of data workflows that are currently executing. You can access your monitoring views from the 'Monitoring' section of the navigation panel. If you have additional workflow management permissions, you can also perform the associated actions from these views.

Once the data workflows that you monitor have completed execution, they appear under 'Completed data workflows', where you can consult their execution history.

## 31.3 Managing work item allocation

If a workflow model defines special allocation management permissions for you or a role that you belong to, you have the ability to manually intervene for work item allocations during the execution of associated data workflows. In this case, you are able to perform one or more of the actions listed below on work items.

Select 'Work items' in the 'Monitoring' section of the navigation pane. The actions that you are able to perform appear in the **Actions** menu of the work item's entry in the table, depending on the current state of the work item.

---

|                   |                                                                                                                          |
|-------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Allocate</b>   | Allocate a work item to a specific user. This action is available for work items in the <i>offered</i> state.            |
| <b>Deallocate</b> | Reset a work item in the <i>allocated</i> state to the <i>offered</i> state.                                             |
| <b>Reallocate</b> | Modify the user to whom a work item is allocated. This action is available for work items in the <i>allocated</i> state. |

---

**See also**

[Work items](#) [p 165]

[Permissions on associated data workflows](#) [p 151]

**Related concepts** [Workflow models](#) [p 130]

## CHAPTER 32

# Administration of data workflows

If you have been given permissions for administration activities associated with data workflows, any relevant publications, active data workflows, and work items, will appear under the entries of the 'Monitoring' section in the navigation panel. From these monitoring views, you can directly perform administrative tasks from the **Actions** menus of the table entries.

**Note**

When a workflow model gives you administrative rights, you automatically have monitoring permissions on all of the relevant aspects of data workflow execution, such as publications, active data workflows, and work items.

This chapter contains the following topics:

1. [Overview of data workflow execution](#)
2. [Data workflow administration actions](#)

## 32.1 Overview of data workflow execution

When a data workflow is launched, a *token* that marks the step currently being executed is created and positioned at the start of the workflow. As each step is completed, this token moves on to the next step as defined in the workflow model on whose publication the data workflow is based.

At any given point during the execution of a data workflow, the token is positioned on one of the following:

- a script task, which is run automatically and requires no user interaction. The script task is completed when the defined actions finish running.
- a user task, which spawns one or more work items to be performed manually by users. Each work item is completed by an explicit 'Accept' or 'Reject' action from a user, and the completion of the umbrella user task is determined according to the task termination criteria defined for the user task in the workflow model.
- a condition, which is evaluated automatically in order to determine the next step in the execution of the data workflow.
- a sub-workflows invocation, which launches associated sub-workflows and waits for the termination of the launched sub-workflows.
- a wait task, which pauses the workflow until a specific event is received.

The token can be in the following states:

- **To execute:** The token is the process of progressing to the next step, based on the workflow model.
- **Executing:** The token is positioned on a script task or a condition that is being processed.
- **User:** The token is positioned on a user task and is awaiting a user action.
- **Waiting for sub-workflows:** The token is positioned on a sub-workflow invocation and is awaiting the termination of all launched sub-workflows.
- **Waiting for event:** The token is positioned on a wait task and is waiting for a specific event to be received.
- **Finished:** The token has reached the end of the data workflow.
- **Error:** An error has occurred.

See also [Workflow management](#) [p 411]

## 32.2 Data workflow administration actions

### *Actions on publications*

#### Disabling a workflow publication

To prevent new data workflows from being launched from a given workflow publication, you can disable it. Select the 'Publications' entry from the navigation pane, then select **Actions > Disable** in the entry for the publication you want to disable.

Once disabled, the publication will no longer appear in the 'Workflow launchers' view of users, but any data workflows already launched that are in progress will continue executing.

#### Note

Once a publication has been disabled, it cannot be re-enabled from the Data Workflows area. Only a user with the built-in repository 'Administrator' role can re-enable a disabled publication from the Administration area, although manually editing technical tables is not generally recommended, as it is important to ensure the integrity of workflow operations.

#### Unpublishing a workflow publication

If a workflow publication is no longer required, you can remove it completely from the views in the Data Workflows area by unpublishing it. To do so,

1. Disable the workflow publication to prevent users from continuing to launch new data workflows from it, as described in [Disabling a workflow publication](#) [p 174].
2. Unpublish the workflow publication by selecting **Actions > Unpublish** from the workflow publication's entry in the same table of publications.

#### Note

When you choose to unpublish a workflow publication, you will be prompted to confirm the termination and cleaning of any data workflows in progress that were launched from this workflow publication, and any associated work items. Any data that is lost as a result of forcefully terminating a data workflow cannot be recovered.

## Actions on data workflows

From the tables of data workflows, it is possible to perform actions from the **Actions** menu in the record of a given data workflow.

### Replaying a step

In the event of an unexpected failure during a step, for example, an access rights issue or unavailable resources, you can "replay" the step as a data workflow administrator. Replaying a step cleans the associated execution environment, including any related work items and sub-workflows, and resets the token to the beginning of the current step.

To replay the current step in a data workflow, select **Actions > Replay the step** from the entry of the workflow in the 'Active workflows' table.

### Terminating and cleaning an active data workflow

In order to stop and clean a data workflow that is currently in progress, select **Actions > Terminate and clean** from the entry of the workflow in the 'Active workflows' table. This will stop the execution of the data workflow and clean the data workflow and all associated work items and sub-workflows.

#### Note

This action is not available on workflows in the 'Executing' state, and on sub-workflows launched from another workflow.

#### Note

Workflow history data is not deleted.

### Forcing termination of an active data workflow

In order to stop a data workflow that is currently in progress, select **Actions > Force termination** from the entry of the workflow in the 'Active workflows' table. This will stop the execution of the data workflow and clean any associated work items and sub-workflows.

#### Note

This action is available for sub-workflows, and for workflows in error blocked on the last step.

#### Note

Workflow history data is not deleted.

### Forcing resumption of a waiting data workflow

In order to resume a data workflow that is currently waiting for an event, select *Actions > Force resumption* from the entry of the workflow in the 'Active workflows' table. This will resume the data workflow. Before doing this action, it is the responsibility of the administrator to update the data context in order to make sure that the data workflow can execute the next steps.

#### Note

This action is only available for workflows in the 'waiting for event' state.

## Cleaning a completed data workflow

When a data workflow has completed its execution, its history is viewable by monitors and administrators of that workflow in the view 'Completed workflows'. To remove the completed workflow, you can perform a clean operation on it. To do so, select **Actions > Clean** from the entry of the workflow in the 'Completed workflows' table.

When cleaned a workflow is no longer visible in the view 'Completed workflows' but its history is still available from the technical administration area.

### Note

This action is not available on sub-workflows launched from another workflow.

See also [Workflow management](#) [p 411]

## Modifying the priority of a data workflow

After a data workflow has been launched, an administrator of the workflow can alter its priority level. Doing so changes the priority of all existing and future work items created by the data workflow. To change the priority level of a data workflow, select **Actions > Modify priority** from the entry of the workflow in the 'Active workflows' table.

See also [Permissions on associated data workflows](#) [p 151]



---

# Data services

---

## CHAPTER 33

# Introduction to data services

This chapter contains the following topics:

1. [Overview](#)
2. [Using the Data Services area user interface](#)

## 33.1 Overview

### ***What is a data service?***

A [data service](#) [p 29] is:

- a standard Web service that interacts with TIBCO EBX®.  
SOAP data services can be dynamically generated based on data models from the 'Data Services' area.
- a REST service that allows interrogating the EBX® repository.  
The built-in RESTful service does not require a service interface, it is self-descriptive through the returned metadata.

They can be used to access some of the features available through the user interface.

See also

[WSDL/SOAP](#) [p 598]

[REST](#) [p 652]


### ***Lineage***

[Lineage](#) [p 30] is used to establish user permission profiles for non-human users, namely data services. When accessing data using WSDL interfaces, data services use the permission profiles established through lineage.

### ***Glossary***

See also [Data services](#) [p 29]

## 33.2 Using the Data Services area user interface



### Note

This area is available only to authorized users in the 'Advanced perspective'.

### Related concepts

[Dataspace](#) [p 88]

[Dataset](#) [p 106]

[Data workflows](#) [p 158]

[Introduction](#) [p 598]



## CHAPTER 34

# Generating data service WSDLs

This chapter contains the following topics:

1. [Generating a WSDL for operations on data](#)
2. [Generating a WSDL for dataspace operations](#)
3. [Generating a WSDL for data workflow operations](#)
4. [Generating a WSDL for lineage](#)
5. [Generating a WSDL for administration](#)
6. [Generating a WSDL to modify the default directory](#)

## 34.1 Generating a WSDL for operations on data

To generate a WSDL for accessing data, select 'Data' in the navigation panel in the **Data Services** area, then follow through the steps of the wizard:

1. Choose whether the WSDL will be for operations at the dataset level or at the table level.
2. Identify the dataspace and dataset on which the operations will be run
3. Select the tables on which the operations are authorized, as well as the operations permitted.
4. Download the generated WSDL file by clicking the button **Download WSDL**.

### ***Operations on datasets***

The following operations can be performed using the WSDL generated for operations at the dataset level:

- Select dataset content for a dataspace or snapshot.
- Get dataset changes between dataspaces or snapshots
- Replication unit refresh

### ***Operations on tables***

The following operations, if selected, can be performed using the WSDL generated for operations at the table level:

- Insert record(s)
- Select record(s)

- Update record(s)
- Delete record(s)
- Count record(s)
- Get changes between dataspace or snapshot
- Get credentials
- Run multiple operations on tables in the dataset

**See also**

[WSDL download from HTTP protocol](#) [p 611]

[Operations generated from a data model](#) [p 617]

## 34.2 Generating a WSDL for dataspace operations

To generate a WSDL for dataspace-level operations, selecting 'Dataspace' in the navigation panel of the **Data Services** area. The generated WSDL is generic to all dataspaces, thus no additional information is required.

Download the generated WSDL file by clicking the button **Download WSDL**.

### ***Operations on dataspaces***

The following operations can be performed using the WSDL generated for operations at the dataspace level:

- Create a dataspace
- Close a dataspace
- Create a snapshot
- Close a snapshot
- Merge a dataspace
- Lock a dataspace
- Unlock a dataspace
- Validate a dataspace or a snapshot
- Validate a dataset

**See also**

[WSDL download from HTTP protocol](#) [p 611]

[Operations on datasets and dataspaces](#) [p 638]

## 34.3 Generating a WSDL for data workflow operations

To generate a WSDL to control data workflows, select 'Data workflow' from the **Data Services** area. The generated WSDL is not specific to any particular workflow publication, thus no additional information is required.

Download the generated WSDL file by clicking the button **Download WSDL**.

## Operations on data workflows

- Start a data workflow
- Resume a data workflow
- End a data workflow

### See also

[WSDL download from HTTP protocol](#) [p 611]

[Operations on data workflows](#) [p 644]

## 34.4 Generating a WSDL for lineage

To generate a WSDL for lineage, select 'Lineage' from the **Data Services** area. It will be based on authorized profiles that have been defined by an administrator in the 'Lineage' section of the **Administration** area.

The operations available for accessing tables are the same as for [WSDL for operations on data](#) [p 181].

Steps for generating the WSDL for lineage are as follows:

1. Select the profile whose permissions will be used. The selected user or role must be authorized for use with lineage by an administrator.
2. Identify the dataspace and dataset on which the operations will be run
3. Select the tables on which the operations are authorized, as well as the operations permitted.
4. Download the generated WSDL file by clicking the button **Download WSDL**.

See also [Lineage](#) [p 178]

## 34.5 Generating a WSDL for administration

*This action is only available to administrators.*

To generate a WSDL for:

- managing the user interface
- getting system information

select 'Administration' from the **Data Services** area.

### Operations for administration

- Close user interface
- Open user interface
- Get system information

### See also

[WSDL download from HTTP protocol](#) [p 611]

[User interface operations](#) [p 648]

[System information operation](#) [p 648]

## 34.6 Generating a WSDL to modify the default directory

*This action is only available to administrators, and only if using the default directory.*

To generate a WSDL to update the default directory, select 'Directory' from the **Data Services** area.

### ***Operations on the default directory***

The operations available for accessing tables are the same as for [WSDL for operations on data](#) [p 181].

#### **See also**

[WSDL download from HTTP protocol](#) [p 611]

[Directory services](#) [p 647]



---

# Reference Manual

---

---

# Integration

---

## CHAPTER 35

# Overview of integration and extension

Several service and component APIs allow you to develop custom extensions for TIBCO EBX® and integrate it with other systems.

This chapter contains the following topics:

1. [Using EBX® as a Web Component](#)
2. [User interface customization](#)
3. [Data services](#)
4. [XML and CSV import/export services](#)
5. [Programmatic services](#)

## 35.1 Using EBX® as a Web Component

It is possible to use EBX® as a user interface web component, by calling it using the HTTP protocol. Such EBX® Web Components can be integrated into any application that is accessible through a [supported web browser](#) [p 308].

A typical use is to integrate EBX® views into an organization's intranet framework. Web Components can also be invoked from the EBX® user interface using [User services](#) [p 187].

See also [Using EBX® as a Web Component](#) [p 191]

## 35.2 User interface customization

### ***User services***

A user service is an extension of EBX® that provides a graphical user interface (GUI) allowing users to access specific or advanced functionalities.

See also [User services overview](#) [p 565]

### ***Custom form layout***

A presentation layer provides the ability to override the default layout of forms in the user interface with highly customized form layouts.

See also [Form layout](#) [p 563]

## Custom widgets

A custom widget is a graphical component developed specifically to customize the look and feel of groups and fields in data models or in programmatically defined schemas.

See also [Custom widgets](#) [p 563]

## Ajax

EBX® supports Ajax asynchronous exchange of data with the server without refreshing the currently displayed page.

See also

[User service Ajax callbacks](#) [p 562]

**Ajax component** `UIAjaxComponent`<sup>API</sup>

## Specifying UI filters on a table

In addition to the default filters and search panes in the user interface, it is possible to specify additional filters dependent on the structure of the table. For that, a specific class must be defined in the definition of the table and must extend `UITableFilter`.

See `UITableFilter`<sup>API</sup> for more information.

## 35.3 Data services

The data services module provides a means for external systems to interact with EBX® using one of following:

- Web Services Description Language (WSDL/SOAP) standard
- Representational state transfer (REST)

See also

[WSDL/SOAP data services](#) [p 598]

[REST data services](#) [p 652]

## 35.4 XML and CSV import/export services

EBX® includes built-in services for importing data from and export data to XML and CSV formats. Imports and exports for XML and CSV can be performed using the user interface, data services, or the Java API.

See also

[XML import and export](#) [p 213]

[CSV import and export](#) [p 219]

## 35.5 Programmatic services

Programmatic services allow executing procedures in a well-defined context, for example in a scheduled task or in a batch.

Some examples of programmatic services include:

- Importing data from an external source,
- Exporting data to multiple systems,
- Data historization, launched by a supervisory system
- Optimizing and refactoring data if EBX® **built-in optimization services** `AdaptationTreeOptimizerSpecAPI` are not sufficient.

**See also** `ProgrammaticServiceAPI`



## CHAPTER 36

# Using TIBCO EBX® as a Web Component

This chapter contains the following topics:

1. [Overview](#)
2. [Integrating EBX® Web Components into applications](#)
3. [Repository element and scope selection](#)
4. [Combined selection](#)
5. [Request specifications](#)
6. [Example calls to an EBX® Web Component](#)

## 36.1 Overview

EBX® can be used as a user interface Web Component, called through the HTTP protocol. An EBX® Web Component can be integrated into any application that is accessible through a supported web browser. This method of access offers the major benefits of EBX®, such as user authentication, data validation, and automatic user interface generation, while additionally providing the ability to focus user navigation on specific elements of the repository.

Typical uses of EBX® Web Components include integrating them into the intranet frameworks of organizations or into applications that manage the assignment of specific tasks to users.

See also [Supported web browsers](#) [p 308]

## 36.2 Integrating EBX® Web Components into applications

A web application that calls an EBX® Web Component can be:

1. A non-Java application, the most basic being a static HTML page.  
In this case, the application must send an HTTP request that follows the EBX® Web Component [request specifications](#) [p 193].
2. A Java application, for example:
  - A Java web application running on the same application server instance as the EBX® repository it targets or on a different application server instance.

- An EBX® [User service](#) [p 187] or a [Custom widget](#) [p 188], in which case, the new session will automatically inherit from the parent EBX® session.

#### Note

In Java, the recommended method for building HTTP requests that call EBX® web components is to use the class `UIHttpManagerComponentAPI` in the API.

## 36.3 Repository element and scope selection

When an EBX® Web Component is called, the user must first be authenticated in the newly instantiated HTTP session. The Web Component then selects a repository element and displays it according to the scope layout parameter defined in the request.

The parameter `firstCallDisplay` may change this automatic display according to its value.

The repository elements that can be selected are as follows:

- Dataspace or snapshot
- Dataset
- Node
- Table or a published view
- Table record

The scope determines how much of the user interface is displayed to the user, thus defining where the user is able to navigate in the session. The default scope that the Web component uses is the smallest possible depending on the entity or service being selected or invoked by the request.

See also [Scope](#) [p 196]

See also [firstCallDisplay](#) [p 196]

It is also possible to select a specific perspective as well as a perspective action.

By default, the selection of the element is done in the context of the perspective of the user if the scope is "full".

See also [Perspective](#) [p 15]

## 36.4 Combined selection

A URL of a Web component can specify a perspective and an action or an entity (dataspace, dataset, etc). Thus, for a Web component that has specified in its URL a perspective and an entity (but no action), if an action of the perspective matches this entity, then this action will be automatically selected.

Otherwise, if no action matches this entity, no action will be selected but the entity is opened regardless.

If an action is specified at the same time than an entity, this last is ignored and the action will be selected.



### ***Specific case***

If the target entity is a record and if an action is on the table that contains this record, then this action will be selected and the record will be opened inside the action.

In the same way, if a workflow work item is targeted by the web component, and if an action on « inbox » exists in the perspective, then this action will be selected and the work item will be opened inside it.

### ***Known limitations***

If the Web component specifies a predicate to filter a table, the perspective action must specify the exact same predicate to be selected.

In the same way, if the perspective action specifies a predicate to filter a table, the Web component must specify the exact same predicate to establish the match.

## **36.5 Request specifications**

### ***Base URL***

In a default deployment, the base URL must be of the following form:

`http://<host>[:<port>]/ebx/`

#### **Note**

The base URL must refer to the servlet `FrontServlet`, defined in the deployment descriptor `/WEB-INF/web.xml` of the web application `ebx.war`.

### ***User authentication and session information parameters***

| Parameter                                                     | Description                                                                                                                                                                                                                                                                                                                                                                    | Required                                           |
|---------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| login and password, or a <i>user directory-specific token</i> | Specifies user authentication properties. If neither a login and password pair nor a user directory-specific token is provided, user will be required to authenticate through the repository login page.<br>See <code>Directory<sup>API</sup></code> for more information.                                                                                                     | No                                                 |
| trackingInfo                                                  | Specifies the tracking information of the new session. Tracking information is logged in history tables. Additionally, it can be used to programmatically restrict access permissions.<br>See <code>AccessRule<sup>API</sup></code> for more information.                                                                                                                      | No                                                 |
| redirect                                                      | The URL to which the user will be redirected at the end of the component session, when they click on the button 'Close'. The close button is always displayed for record selections, but whether or not it is displayed for all other cases must be specified using the parameter <code>closeButton</code> .<br>For more information, see <a href="#">Exit policy</a> (p 389). | No                                                 |
| locale                                                        | Specifies the locale to use. Value is either <code>en-US</code> or <code>fr-FR</code> .                                                                                                                                                                                                                                                                                        | No, default is the locale registered for the user. |

## Entity and service selection parameters

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Required                                                           |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------|
| branch              | Selects the specified dataspace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | No                                                                 |
| version             | Selects the specified snapshot.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | No                                                                 |
| instance            | Selects the specified dataset. The value must be the reference of a dataset that exists in the selected dataspace or snapshot.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Only if xpath or viewPublication is specified.                     |
| viewPublication     | <p>Specifies the publication name of the tabular or hierarchical view to apply to the selected content.</p> <p>This publication name is the one declared during the publication of the view. It can be found in the 'Administration' area under <i>Views configuration &gt; Views</i>.</p> <p>All settings of the view, that is, its filters, sort order, and displayed columns, are applied to the result. A dataspace and a dataset must be selected in order for this view to be applied. The target table selection is not necessary, as it can be automatically determined based on the definition of the view. This parameter can be combined with the predicate specified in the xpath parameter as a logical 'AND' operation.</p> | No                                                                 |
| xpath               | <p>Specifies a node selection in the dataset.</p> <p>Value may be a valid absolute path located in the selected dataset. The notation must conform to a simplified XPath, with abbreviated syntax.</p> <p>It can also be a predicate surrounded by "[" and "]" if a table can be automatically selected using other Web Component parameters (for example, viewPublication or workflowView).</p> <p>For XPath syntax, see <a href="#">XPath supported syntax</a> [p 225]</p> <p>See <code>UIHttpManagerComponent.setPredicate<sup>API</sup></code> for more information.</p>                                                                                                                                                              | No                                                                 |
| service             | <p>Specifies the service to access.</p> <p>For more information on built-in User services, see <a href="#">Built-in services</a> [p 199].</p> <p>In the Java API, see <code>ServiceKey<sup>API</sup></code> for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | No                                                                 |
| workflowView        | <p>Specifies the workflow section to be selected.</p> <p>See <code>workflowView<sup>API</sup></code> for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | No.                                                                |
| perspectiveName     | <p>Specifies the name of the perspective to be selected.</p> <p>If this parameter is specified, the scope parameter can have only two values: full and data.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Only if perspectiveActionId or perspectiveActionName is specified. |
| perspectiveActionId | <p><b>Deprecated.</b> Please consider using perspectiveActionName instead.</p> <p>Specifies the identifier of the perspective action to be selected.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | No.                                                                |

| Parameter             | Description                                                         | Required |
|-----------------------|---------------------------------------------------------------------|----------|
| perspectiveActionName | Specifies the unique name of the perspective action to be selected. | No.      |

## Layout parameters

| Parameter        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Required                                                                                    |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| scope            | Specifies the scope to be used by the web component. Value can be full, data, dataspace, dataset or node.<br>See <code>UIHttpManagerComponent.Scope<sup>API</sup></code> for more information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | No, default will be computed to be the smallest possible according to the target selection. |
| firstCallDisplay | Specifies which display must be used instead of the one determined by the combination of selection and scope parameter.<br>Possible values are: <ul style="list-style-type: none"> <li>• auto: The display is automatically set according to the selection.</li> <li>• view: Forces the display of the tabular view or of the hierarchical view.</li> <li>• record: If the predicate has at least one record, forces the display of the first record in the list.</li> </ul> For example,<br><code>firstCallDisplay=view</code><br><code>firstCallDisplay=view:hierarchyExpanded</code><br><code>firstCallDisplay=record</code><br><code>firstCallDisplay=record:{predicate}</code><br>See <code>UIHttpManagerComponent.setFirstCallDisplay<sup>API</sup></code> for more information.<br>See <code>UIHttpManagerComponent.setFirstCallDisplayHierarchyExpanded<sup>API</sup></code> for more information.<br>See <code>UIHttpManagerComponent.setFirstCallDisplayRecord<sup>API</sup></code> for more information. | No, default will be computed according to the target selection.                             |
| closeButton      | Specifies how to display the session close button. Value can be logout or cross.<br>See <code>UIHttpManagerComponent.CloseButtonSpec<sup>API</sup></code> for more information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | No. If scope is not full, no close button will be displayed by default.                     |
| dataSetFeatures  | Specifies which features to display in a UI service at the dataset level or a form outside of a table.<br>These options pertain only to features in the workspace. It is recommended to use this property with the smallest scope possible, namely dataset or node.<br>Syntax:<br><code>&lt;prefix&gt; ":" &lt;feature&gt; [ ", " &lt;feature&gt; ]*</code><br>where <ul style="list-style-type: none"> <li>• <code>&lt;prefix&gt;</code> is hide or show,</li> <li>• <code>&lt;feature&gt;</code> is services, title, save, or revert.</li> </ul> For example,<br><code>hide:title</code><br><code>show:save, revert</code><br>See <code>UIHttpManagerComponent.DataSetFeatures<sup>API</sup></code> for more information.                                                                                                                                                                                                                                                                                         | No.                                                                                         |
| viewFeatures     | Specifies which features to display in a tabular or a hierarchy view (at the table level).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | No.                                                                                         |

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Required                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
|                | <p>These options pertain only to features in the workspace. It is recommended to use this property with the smallest scope possible, namely dataset or node.</p> <p>Syntax:</p> <pre>&lt;prefix&gt; ":" &lt;feature&gt; [ ", " &lt;feature&gt;]*</pre> <p>where</p> <ul style="list-style-type: none"> <li>• &lt;prefix&gt; is hide or show,</li> <li>• &lt;feature&gt; is create, views, selection, filters, services, refresh, title, or breadcrumb.</li> </ul> <p>For example,</p> <pre>hide:title,selection</pre> <pre>show:service,title,breadcrumb</pre> <p>See <code>UIHttpManagerComponent.ViewFeatures<sup>API</sup></code> for more information.</p>                                                                 |                                                                                 |
| recordFeatures | <p>Specifies which features must be displayed in a form at the record level.</p> <p>These options pertain only to features in the workspace. It is recommended to use this property with the smallest scope possible, namely dataset or node.</p> <p>Syntax:</p> <pre>&lt;prefix&gt; ":" &lt;feature&gt; [ ", " &lt;feature&gt;]*</pre> <p>where</p> <ul style="list-style-type: none"> <li>• &lt;prefix&gt; is hide or show,</li> <li>• &lt;feature&gt; is services, title, breadcrumb, save, saveAndClose, close, or revert.</li> </ul> <p>For example,</p> <pre>hide:title</pre> <pre>show:save,saveAndClose,revert</pre> <p>See <code>UIHttpManagerComponent.RecordFeatures<sup>API</sup></code> for more information.</p> | No.                                                                             |
| pageSize       | Specifies the number of records that will be displayed per page in a table view (either tabular or hierarchical).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | No.                                                                             |
| startWorkItem  | <p>Specifies a work item must be automatically taken and started. Value can be true or false.</p> <p>See <code>ServiceKey.WORKFLOW<sup>API</sup></code> for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | No. Default value is false, where the target work item state remains unchanged. |

## 36.6 Example calls to an EBX® Web Component

Minimal URI:

```
http://localhost:8080/ebx/
```

Logs in as the user 'admin' and selects the 'Reference' dataspace:

```
http://localhost:8080/ebx/?login=admin&password=admin&branch=Reference
```

Selects the 'Reference' dataspace and accesses the built-in validation service:

```
http://localhost:8080/ebx/?login=admin&password=admin&branch=Reference&service=@validation
```

Selects the roles table in the default directory:

```
http://localhost:8080/ebx/?login=admin&password=admin&branch=ebx-directory&instance=ebx-directory&xpath=/directory/roles
```

Selects the record 'admin' in the default directory:

```
http://localhost:8080/ebx/?login=admin&password=admin&branch=ebx-  
directory&instance=ebx-directory&xpath=/directory/users[./login="admin"]
```

Accesses the interface for creating a new user in the default directory:

```
http://localhost:8080/ebx/?login=admin&password=admin&branch=ebx-  
directory&instance=ebx-directory&xpath=/directory/users&service=@creation
```

Compares the record 'admin' in the default directory with the record 'jSmith':

Compares the record 'R1' in the dataset 'instanceId' in the dataspace 'Reference' with the record 'R0':

```
http://localhost:8080/ebx/?login=admin&password=admin&branch=ebx-  
directory&instance=ebx-directory&xpath=/directory/users[./  
login="admin"]&service=@compare&compare.branch=ebx-directory&compare.instance=ebx-  
directory&compare.xpath=/directory/users[./login="jSmith"]
```

## CHAPTER 37

# Built-in user services

EBX® includes a number of built-in user services. Built-in user services can be used:

- [when defining workflow model tasks](#) [p 136]
- [when defining perspective action menu items](#) [p 16]
- [as extended user services when used with service extensions](#) [p 589]
- [when using EBX® as a Web Component](#) [p 191]

This reference page describes the built-in user services and their parameters.

This chapter contains the following topics:

1. [Access data \(default service\)](#)
2. [Create a new record](#)
3. [Duplicate a record](#)
4. [Export data to an XML file](#)
5. [Export data to a CSV file](#)
6. [Import data from an XML file](#)
7. [Import data from a CSV file](#)
8. [Access a dataspace](#)
9. [Validate a dataspace, a snapshot or a dataset](#)
10. [Merge a dataspace](#)
11. [Access the dataspace merge view](#)
12. [Compare contents](#)
13. [Data workflows](#)

## 37.1 Access data (default service)

By default, workflows automatically consider this service as complete. That is, the 'Accept' button is always available.

This is the default service used if no service is specified.

## Input parameters

| Parameter           | Label                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch              | Dataspace               | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                                                                                                                                                                         |
| disableAutoComplete | Disable Accept at start | By default, the interaction associated with this service is directly considered as complete. Therefore, the Accept button is automatically displayed at the opening of the work item. This parameter is useful to disable this behavior. If the value is 'true', the developer will be in charge of completing the interaction by using SessionInteraction in a user service or a trigger, for example. The default value is 'false'. Perspectives do not use this parameter. |
| firstCallDisplay    | First call display mode | Defines the display mode that must be used when displaying a filtered table or a record upon first call. Default (value = 'auto'): the display is automatically set according to the selection. View (value = 'view'): forces the display of the tabular view or of the hierarchical view. Record (value = 'record'): if the predicate has at least one record, forces the display of the record form.                                                                        |
| instance            | Dataset                 | The value must be the reference of a dataset that exists in the selected dataspace.                                                                                                                                                                                                                                                                                                                                                                                           |
| scope               | Scope                   | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service.                                                                                                                                                      |
| trackingInfo        | Tracking information    | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                                                                                                                                                                           |
| version             | Snapshot                | The identifier of the specified snapshot - A dataspace is required for this service.                                                                                                                                                                                                                                                                                                                                                                                          |
| viewPublication     | View                    | The publication name of the view to display. The view must be configured for the selected table.                                                                                                                                                                                                                                                                                                                                                                              |
| xpath               | Dataset node (XPath)    | The value must be a valid absolute location path in the selected dataset. The notation must conform to a simplified XPath, in its abbreviated syntax.                                                                                                                                                                                                                                                                                                                         |



## 37.2 Create a new record

For a workflow, the creation service is considered complete when the first successful submit is performed (record has been created). If this service is called whereas it is already complete, the created record is displayed in update or read-only mode (depending on the user rights).

Service name parameter: `service=@creation`

### *Input parameters*

| Parameter    | Label                 | Description                                                                                                                                                                                                                                                                                                              |
|--------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace             | The identifier of the specified dataspace - This field is required for this service.                                                                                                                                                                                                                                     |
| instance     | Dataset               | The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.                                                                                                                                                                                            |
| scope        | Scope                 | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information  | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| xpath        | Dataset table (XPath) | The value must be a valid absolute location path of a table in the selected dataset. The notation must conform to a simplified XPath, in its abbreviated syntax - This field is required for this service.                                                                                                               |

### *Output parameters*

| Parameter | Label          | Description                               |
|-----------|----------------|-------------------------------------------|
| created   | Created record | Contains the XPath of the created record. |

## 37.3 Duplicate a record

For a workflow, the duplicate service is considered complete when the first successful submit is performed (record has been created). If this service is called whereas it is already complete, the created record is displayed in update or read-only mode (depending on the user rights).

Service name parameter: `service=@duplicate`

## ***Input parameters***

| Parameter    | Label                       | Description                                                                                                                                                                                                                                                                                                              |
|--------------|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace                   | The identifier of the specified dataspace - This field is required for this service.                                                                                                                                                                                                                                     |
| instance     | Dataset                     | The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.                                                                                                                                                                                            |
| scope        | Scope                       | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information        | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| xpath        | Record to duplicate (XPath) | The value must be a valid absolute location path of an existing record. The notation must conform to a simplified XPath, in its abbreviated syntax - This field is required for this service.                                                                                                                            |

## ***Output parameters***

| Parameter | Label          | Description                               |
|-----------|----------------|-------------------------------------------|
| created   | Created record | Contains the XPath of the created record. |

## **37.4 Export data to an XML file**

The exportToXML service is considered complete when export is done and file downloaded.

Service name parameter: service=@exportToXML

## *Input parameters*

| Parameter    | Label                           | Description                                                                                                                                                                                                                                                                                                              |
|--------------|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace                       | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                    |
| instance     | Dataset                         | The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.                                                                                                                                                                                            |
| scope        | Scope                           | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information            | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| version      | Snapshot                        | The identifier of the specified snapshot - A dataspace is required for this service.                                                                                                                                                                                                                                     |
| xpath        | Dataset table to export (XPath) | The value must be a valid absolute location path of a table in the selected dataset. The notation must conform to a simplified XPath, in its abbreviated syntax - This field is required for this service.                                                                                                               |

## 37.5 Export data to a CSV file

Workflows consider the exportToCSV service as complete when export is done and file downloaded.

Service name parameter: `service=@exportToCSV`

## ***Input parameters***

| Parameter    | Label                           | Description                                                                                                                                                                                                                                                                                                              |
|--------------|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace                       | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                    |
| instance     | Dataset                         | The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.                                                                                                                                                                                            |
| scope        | Scope                           | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information            | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| version      | Snapshot                        | The identifier of the specified snapshot - A dataspace is required for this service.                                                                                                                                                                                                                                     |
| xpath        | Dataset table to export (XPath) | The value must be a valid absolute location path of a table in the selected dataset. The notation must conform to a simplified XPath, in its abbreviated syntax - This field is required for this service.                                                                                                               |

## **37.6 Import data from an XML file**

Workflows consider the importFromXML service as complete when import is performed.

Service name parameter: service=@importFromXML

## ***Input parameters***

| Parameter    | Label                           | Description                                                                                                                                                                                                                                                                                                              |
|--------------|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace                       | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                    |
| instance     | Dataset                         | The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.                                                                                                                                                                                            |
| scope        | Scope                           | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information            | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| xpath        | Dataset table to import (XPath) | The value must be a valid absolute location path of a table in the selected dataset. The notation must conform to a simplified XPath, in its abbreviated syntax - This field is required for this service.                                                                                                               |

## **37.7 Import data from a CSV file**

Workflows consider the importFromCSV service as complete when import is performed.

Service name parameter: service=@importFromCSV

## ***Input parameters***

| Parameter    | Label                           | Description                                                                                                                                                                                                                                                                                                              |
|--------------|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace                       | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                    |
| instance     | Dataset                         | The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.                                                                                                                                                                                            |
| scope        | Scope                           | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information            | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| xpath        | Dataset table to import (XPath) | The value must be a valid absolute location path of a table in the selected dataset. The notation must conform to a simplified XPath, in its abbreviated syntax - This field is required for this service.                                                                                                               |

## **37.8 Access a dataspace**

A workflow automatically considers that the dataspace selection service is complete.

Service name parameter: `service=@selectDataSpace`

## Input parameters

| Parameter    | Label                | Description                                                                                                                                                                                                                                                                                                              |
|--------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace            | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                    |
| scope        | Scope                | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| version      | Snapshot             | The identifier of the specified snapshot - A dataspace is required for this service.                                                                                                                                                                                                                                     |

## 37.9 Validate a dataspace, a snapshot or a dataset

Workflows automatically consider the validation service as complete.

Service name parameter: service=@validation

## Input parameters

| Parameter    | Label                | Description                                                                                                                                                                                                                                                                                                              |
|--------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace            | The identifier of the specified dataspace - A dataspace or snapshot is required for this service.                                                                                                                                                                                                                        |
| instance     | Dataset              | The value must be the reference of a dataset that exists in the selected dataspace.                                                                                                                                                                                                                                      |
| scope        | Scope                | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| version      | Snapshot             | The identifier of the specified snapshot - A dataspace or snapshot is required for this service.                                                                                                                                                                                                                         |

## Output parameters

| Parameter  | Label              | Description                                              |
|------------|--------------------|----------------------------------------------------------|
| hasError   | Found errors       | Contains 'true' if validation has produced errors.       |
| hasFatal   | Found fatal errors | Contains 'true' if validation has produced fatal errors. |
| hasInfo    | Found informations | Contains 'true' if validation has produced informations. |
| hasWarning | Found warnings     | Contains 'true' if validation has produced warnings.     |

## 37.10 Merge a dataspace

Workflows consider the merge service as complete when merger is performed and dataspace is closed.

Service name parameter: `service=@merge`

## Input parameters

| Parameter    | Label                | Description                                                                                                                                                                                                                                                                                                              |
|--------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace            | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                    |
| scope        | Scope                | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |

## Output parameters

| Parameter   | Label         | Description                                                                                                                                                |
|-------------|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| mergeResult | Merge success | Contains 'true' if merge succeeded, otherwise 'false'.                                                                                                     |
| mergeState  | Merge state   | Contains the return code of the merge. It is strongly recommended to parse this value by using the InteractionMergeState UIHttpManagerComponentReturnCode. |



## 37.11 Access the dataspace merge view

The merge.view service is automatically considered complete.

Service name parameter: service=@merge.view

### *Input parameters*

| Parameter    | Label                | Description                                                                                                                                                                                                                                                                                                              |
|--------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch       | Dataspace            | The identifier of the specified dataspace - A dataspace is required for this service.                                                                                                                                                                                                                                    |
| scope        | Scope                | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo | Tracking information | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |

## 37.12 Compare contents

Workflows automatically consider the compare service as complete.

Service name parameter: service=@compare

## Input parameters

| Parameter        | Label                              | Description                                                                                                                                                                                                                                                                                                              |
|------------------|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch           | Dataspace                          | The identifier of the specified dataspace - A dataspace or snapshot and a dataspace or snapshot to compare to are required for this service.                                                                                                                                                                             |
| compare.branch   | Dataspace to compare               | The identifier of the dataspace to compare - A dataspace or snapshot and a dataspace or snapshot to compare to are required for this service.                                                                                                                                                                            |
| compare.filter   | Comparison filter                  | To ignore inheritance and computed values fields in the comparison (disable resolved mode), the filter "persistedValuesOnly" must be specified. By default, when no filter is defined, the comparison uses resolved mode.                                                                                                |
| compare.instance | Dataset to compare                 | The value must be the reference of a dataset that exists in the selected dataspace to compare.                                                                                                                                                                                                                           |
| compare.version  | Snapshot to compare                | The identifier of the snapshot to compare - A dataspace or snapshot and a dataspace or snapshot to compare to are required for this service.                                                                                                                                                                             |
| compare.xpath    | Table or record to compare (XPath) | The value must be a valid absolute location path of a table or a record in the selected dataset to compare. The notation must conform to a simplified XPath, in its abbreviated syntax.                                                                                                                                  |
| instance         | Dataset                            | The value must be the reference of a dataset that exists in the selected dataspace.                                                                                                                                                                                                                                      |
| scope            | Scope                              | Defines the scope of the user navigation for this service, namely, the entities that the user is able to select during their session. If left blank, the default value will be used. For perspectives, the default value is always 'node'. For workflows, the default value depends on the selected entities or service. |
| trackingInfo     | Tracking information               | Tracking information is logged into 'history' logs. It may also be used for any other purpose like access control or additional export information.                                                                                                                                                                      |
| version          | Snapshot                           | The identifier of the specified snapshot - A dataspace or snapshot and a dataspace or snapshot to compare to are required for this service.                                                                                                                                                                              |
| xpath            | Table or record (XPath)            | The value must be a valid absolute location path of a table or a record in the selected dataset. The notation must                                                                                                                                                                                                       |

| Parameter | Label | Description                                               |
|-----------|-------|-----------------------------------------------------------|
|           |       | conform to a simplified XPath, in its abbreviated syntax. |

## 37.13 Data workflows

This service provides access to the data workflows user interfaces.

Service name parameter: `service=@workflow`

### Note

This service is for perspectives only.

### *Input parameters*

| Parameter       | Label            | Description                                                                                                                                                                               |
|-----------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scope           | Scope            | Defines the scope of the user navigation for this service.                                                                                                                                |
| viewPublication | View publication | Defines the publication name of the view to apply for this service.                                                                                                                       |
| workflowView    | Workflow view    | Specifies the workflow view type. Value can be one of the following: "inbox", "launcher", "monitoringPublications", "monitoringWorkflows", "monitoringWorkItems" or "completedWorkflows". |
| xpath           | Filter (XPath)   | An optional filter. The syntax should conform to an XPath predicate surrounded by "[" and "]".                                                                                            |



## CHAPTER 38

# XML import and export

This chapter contains the following topics:

1. [Introduction](#)
2. [Imports](#)
3. [Exports](#)
4. [Handling of field values](#)
5. [Known limitations](#)

## 38.1 Introduction

XML imports and exports can be performed on tables through the user interface using the **Actions** menu in the workspace.

Both imports and exports are performed in the context of a dataset.

Imports and exports can also be done programmatically.

Default import and export option values can be set in the Administration area, under *User interface > Graphical interface configuration > Default option values > Import/Export*.

## 38.2 Imports

### Attention

Imported XML documents must be encoded in UTF-8 and its structure must conform to the underlying data model of the target dataset.

## ***Import mode***

When importing an XML file, you must specify one of the following import modes, which will dictate how the import procedure handles the source records.

|                                       |                                                                                                                                                                                                                                                                                     |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Insert mode</b>                    | Only record creations are allowed. If a record exists in the target table with the same primary key as the source record, an error is returned and the whole import operation is cancelled.                                                                                         |
| <b>Update mode</b>                    | Only modifications of existing records are allowed. If no record exists in the target table with the same primary key as the source record, an error is returned and the whole import operation is cancelled.                                                                       |
| <b>Update or insert mode</b>          | If a record with the same primary key as the source record already exists in the target table, that record is updated. Otherwise, a new record is created.                                                                                                                          |
| <b>Replace (synchronization) mode</b> | If a record with the same primary key as the source record already exists in the target table, that record is updated. Otherwise, a new record is created. If a record exists in the target table but is not present in the source XML file, that record is deleted from the table. |

## ***Insert and update operations***

The mode *'by delta'* allows ignoring data model elements that are missing from the source XML document. This mode can be enabled through data services or the Java API. The following table summarizes the behavior of insert and update operations when elements are not present in the source document.

See the data services operations [update](#) [p 625] and [insert](#) [p 627], as well as `ImportSpec.setByDelta`<sup>API</sup> in the Java API for more information.

| State in source XML document                                        | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Element does not exist in the source document                       | <p><b>If 'by delta' mode is disabled (default):</b></p> <p>Target field value is set to one of the following:</p> <ul style="list-style-type: none"> <li>• If the element defines a default value, the target field value is set to that default value.</li> <li>• If the element is of a type other than a string or list, the target field value is set to null.</li> <li>• If the element is an aggregated list, the target field value is set to an empty list.</li> <li>• If the element is a string that distinguishes null from an empty string, the target field value is set to null. If it is a string that does not distinguish between the two, an empty string.</li> <li>• If the element (simple or complex) is hidden in data services, the target value is not changed.</li> </ul> <p>See also <a href="#">Hiding a field in Data Services</a> [p 542]</p> <p><b>Note:</b> The user performing the import must have the permissions necessary to create or change the target field value. Otherwise, the value will remain unchanged.</p> <p><b>If 'by delta' mode has been enabled through data services or the Java API:</b></p> <ul style="list-style-type: none"> <li>• For the update operation, the field value remains unchanged.</li> <li>• For the insert operation, the behavior is the same as when byDelta mode is disabled.</li> </ul> |
| Element exists but is empty (for example, <fieldA/>)                | <ul style="list-style-type: none"> <li>• For nodes of type xs:string (or one of its sub-types), the target field's value is set to null if it distinguishes null from an empty string. Otherwise, the value is set to empty string.</li> <li>• For non-xs:string type nodes, an exception is thrown in conformance with XML Schema.</li> </ul> <p>See also <a href="#">TIBCO EBX® whitespace management for data types</a> [p 527]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Element is present and null (for example, <fieldA xsi:nil="true"/>) | <p>The target field is always set to null except for lists, for which it is not supported.</p> <p>In order to use the <code>xsi:nil="true"</code> attribute, you must import the namespace declaration <code>xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Set missing values as null

When updating existing records, if a node is missing or empty in the XML file: if this option is "yes", it will be considered as null. If this option is "no", it will not be modified.

### Ignore extra columns

It may happen that the XML document contains elements that do not exist in the target data model. By default, in this case, the import procedure will fail. It is possible, however, to allow users to launch import procedures that will ignore the extra columns defined in the XML files. This can be done in the configuration parameters of the import wizard for XML. The default value of this parameter can be configured in the 'User interface' configuration under the 'Administration' area.

### *Optimistic locking*

If the technical attribute `ebxd:lastTime` exists in the source XML file, the import mechanism performs a verification to prevent an update operation on a record that may have changed since the last read. In order to use the `ebxd:lastTime` attribute, you must import the namespace declaration `xmlns:ebxd="urn:ebx-schemas:deployment_1.0"`. The timestamp associated with the current record will be compared to this timestamp. If they are different, the update is rejected.

## 38.3 Exports

### Note

Exported XML documents are always encoded in UTF-8.

When exporting to XML, if the table has filters applied, only the records that correspond to the filter are included in the exported file.



The XML export options are as follows:

|                               |                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Download file name</b>     | Specifies the name of the XML file to be exported. This field is pre-populated with the name of the table from which the records are being exported.                                                                                                                                                                                                                                           |
| <b>User-friendly mode</b>     | Specifies whether exported values will be represented in a user-friendly way, or in the standard XML raw format. For example, in user-friendly mode, dates and numbers are formatted according to the user's locale, and foreign keys and enumerated values display their associated labels.<br><b>Note:</b> If this option is selected, the exported file will not be able to be re-imported. |
| <b>Include technical data</b> | Specifies whether internal technical data will be included in the export.<br><b>Note:</b> If this option is selected, the exported file will not be able to be re-imported.                                                                                                                                                                                                                    |
| <b>Is indented</b>            | Specifies whether the file should be indented to improve its readability by a human.                                                                                                                                                                                                                                                                                                           |
| <b>Omit XML comment</b>       | Specifies whether the generated XML comment that describes the location of data and the date of the export should be omitted from the file.                                                                                                                                                                                                                                                    |

## 38.4 Handling of field values

### *Date, time & dateTime format*

The following date and time formats are supported:

| Type        | Format                                         | Example             |
|-------------|------------------------------------------------|---------------------|
| xs:date     | yyyy-MM-dd                                     | 2007-12-31          |
| xs:time     | HH:mm:ss or HH:mm:ss.SSS                       | 11:55:00            |
| xs:dateTime | yyyy-MM-ddTHH:mm:ss or yyyy-MM-ddTHH:mm:ss.SSS | 2007-12-31T11:55:00 |

## 38.5 Known limitations

### ***Association fields***

The XML import and export services do not support association values.

Exporting such fields will not cause any error, however, no value will be exported.

Importing such fields will cause an error, and the import procedure will be aborted.

### ***Selection nodes***

The XML import and export services do not support selection values.

Exporting such fields will not cause any error, however, no value will be exported.

Importing such fields will cause an error, and the import procedure will be aborted.

## CHAPTER 39

# CSV import and export

This chapter contains the following topics:

1. [Introduction](#)
2. [Exports](#)
3. [Imports](#)
4. [Handling of field values](#)
5. [Known limitations](#)

## 39.1 Introduction

CSV imports and exports can be performed on tables through the user interface using the **Actions** menu in the workspace.

Both imports and exports are performed in the context of a dataset.

Imports and exports can also be done programmatically.

Default import and export option values can be set in the Administration area, under *User interface > Graphical interface configuration > Default option values > Import/Export*.

See also [Default option values](#) [p 393]

## 39.2 Exports

When exporting to CSV, if the table has filters applied, only the records that correspond to the filter are included in the exported file.

The CSV export options are as follows:

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Download file name</b>     | Specifies the name of the CSV file to be exported. This field is pre-populated with the name of the table from which the records are being exported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>File encoding</b>          | Specifies the character encoding to use for the exported file. The default is UTF-8.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Enable inheritance</b>     | <p>In order to consider the <a href="#">inheritance</a> [p 25] during a CSV export, the option has to be defined in the model.</p> <p>For more information on inheritance, see <a href="#">Inheritance and value resolution</a> [p 268].</p> <p>Specifies if inheritance will be taken into account during a CSV export.</p> <p>If inheritance is enabled, resolved values of fields are exported with the technical data that define the possible inheritance mode of the record or the field.</p> <p>If inheritance is disabled, resolved values of fields are exported and occulted records are ignored.</p> <p>By default, this option is disabled.</p> <p><b>Note:</b> Inheritance is always ignored, if the table dataset has no parent or if the table has no inherited field.</p> |
| <b>User-friendly mode</b>     | <p>Specifies whether exported values will be represented in a user-friendly way, or in a raw format. For example, in user-friendly mode, dates and numbers are formatted according to the user's locale, and foreign keys and enumerated values display their associated labels.</p> <p><b>Note:</b> If this option is selected, the exported file will not be able to be re-imported.</p>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Include technical data</b> | <p>Specifies whether internal technical data will be included in the export.</p> <p><b>Note:</b> If this option is selected, the exported file will not be able to be re-imported.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Column header</b>          | <p>Specifies whether or not to include column headers in the CSV file.</p> <ul style="list-style-type: none"> <li>• <b>No header</b></li> <li>• <b>Label:</b> For each column in the spreadsheet, the CSV displays its label. Each label is localized according to the locale preference of the current session. If no user-</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

friendly label is defined for a node, the technical name of the node is used.

- **XPath:** For each column in the spreadsheet, the CSV displays the path to the node in the table.

---

**Field separator**

Specifies the field separator to use for exports. The default separator is comma, it can be modified under *Administration > User interface*.

---

**List separator**

Specifies the separator to use for values lists. The default separator is line return, it can be modified under *Administration > User interface*.

---

Programmatic CSV exports are performed using the classes `ExportSpecAPI` and `ExportImportCSVSpecAPI` in the Java API.

## 39.3 Imports

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Download file name</b> | Specifies the name of the CSV file to be imported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Import mode</b>        | <p>When importing a CSV file, you must specify one of the following import modes, which will control the integrity of operations between the source and the target table.</p> <ul style="list-style-type: none"> <li>• <b>Insert mode:</b> Only record creation is allowed. If a record exists in the target table with the same primary key as the source record, an error is returned and the whole import operation is cancelled.</li> <li>• <b>Update mode:</b> Only modifications of existing records are allowed. If no record exists in the target table with the same primary key as the source record, an error is returned and the whole import operation is cancelled.</li> <li>• <b>Update or insert mode:</b> If a record with the same primary key as the source record already exists in the target table, that record is updated. Otherwise, a new record is created.</li> <li>• <b>Replace (synchronization) mode:</b> If a record with the same primary key as the source record already exists in the target table, that record is updated. Otherwise, a new record is created. If a record exists in the target table but is not present in the source XML file, that record is deleted from the table.</li> </ul> |
| <b>File encoding</b>      | Specifies the character encoding to use for the exported file. The default is UTF-8.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Enable inheritance</b> | <p>In order to consider the <a href="#">inheritance</a> [p 25] during a CSV import, the option has to be defined in the model.</p> <p>For more information on inheritance, see <a href="#">Inheritance and value resolution</a> [p 268] and <code>ExportImportCSVSpec.setInheritanceEnabled<sup>API</sup></code>.</p> <p>Specifies whether the inheritance will be taken into account during a CSV import. If technical data in the CSV file define an inherit mode, corresponding fields or records are forced to be inherited. If technical data define an occult mode, corresponding records are forced to be occulted. Otherwise, fields are overwritten with values read from the CSV file. By default, this option is disabled.</p> <p><b>Note:</b> Inheritance is always ignored if the dataset of the table has no parent or if the table has no inherited field.</p>                                                                                                                                                                                                                                                                                                                                                          |

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Column header</b>   | <p>Specifies whether or not to include column headers in the CSV file.</p> <ul style="list-style-type: none"> <li>• <b>No header</b></li> <li>• <b>Label:</b> For each column in the spreadsheet, the CSV displays its label. Each label is localized according to the locale preference of the current session. If no user-friendly label is defined for a node, the technical name of the node is used.</li> <li>• <b>XPath:</b> For each column in the spreadsheet, the CSV displays the path to the node in the table.</li> </ul> |
| <b>Field separator</b> | <p>Specifies the field separator to use for exports. The default separator is comma, it can be modified under <i>Administration &gt; User interface</i>.</p>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List separator</b>  | <p>Specifies the separator to use for values lists. The default separator is line return, it can be modified under <i>Administration &gt; User interface</i>.</p>                                                                                                                                                                                                                                                                                                                                                                     |

Programmatic CSV imports are performed using the classes `ImportSpecAPI` and `ExportImportCSVSpecAPI` in the Java API.

## 39.4 Handling of field values

### ***Aggregated lists***

The CSV import and export services support multi-valued fields, namely aggregated lists. This is only supported for simple typed lists, such as lists of `string`, `date`, or `int`, and for foreign keys. If a table reference is linked to a composite primary key, each item in the list is a formatted string, for example, "true|99". Aggregated lists of groups are not exported.

At export, the items in the list are separated using line separators. In cases where the exported field already contains a line separator, for example in an `osd:html` or an `osd:text`, the code `_crnl_` is inserted in place of the field value's line separators. The same formatting is expected at import, with the whole field value surrounded by quotes.

### ***Hidden fields***

Hidden fields are exported as `ebx-csv:hidden` strings. An imported hidden string will not modify a field's content.

### ***'Null' value for strings***

Using CSV import and export services, a string with a value set to `null` is exported as an empty string. Therefore, a round trip export-import procedure will end up replacing `null` string values with empty strings.

Using programmatic services, the specific value `ebx-csv:nil` can be assigned to strings with values set to `null`. If this is done, the `null` string values will not be replaced by empty strings during round

trip export-import procedures. See `ExportImportCSVSpec.setNullStringEncodedAPI` in the Java API for more information.

### ***Date, time & dateTime format***

The following date and time formats are supported:

| Type        | Format                                         | Example             |
|-------------|------------------------------------------------|---------------------|
| xs:date     | yyyy-MM-dd                                     | 2007-12-31          |
| xs:time     | HH:mm:ss or HH:mm:ss.SSS                       | 11:55:00            |
| xs:dateTime | yyyy-MM-ddTHH:mm:ss or yyyy-MM-ddTHH:mm:ss.SSS | 2007-12-31T11:55:00 |

## 39.5 Known limitations

### ***Aggregated lists of groups***

The CSV import and export services do not support multi-valued groups, that is, aggregated lists of complex type elements. Exporting such nodes will not cause any error, however, no value will be exported.

### ***Terminal groups***

In a CSV file, it is impossible to differentiate a created terminal group that contains only empty fields from a non-created one.

As a consequence, some differences may appear during comparison after performing an export followed by an import. To ensure the symmetry of import and export, use XML import and export instead. See [XML import and export](#) [p 213].

### ***Column label headers***

If two columns share the same label header, an export of the table can be performed successfully, but exported data cannot later be re-imported.

### ***Association fields***

The CSV import and export services do not support association values, i.e. the associated records.

Exporting such fields will not cause any error, however, no value will be exported.

Importing such fields will cause an error and the import procedure will be aborted.

### ***Selection nodes***

The CSV import and export services do not support selection values, i.e. the selected records.

Exporting such fields will not cause any error, however, no value will be exported.

Importing such fields will cause an error and the import procedure will be aborted.



## CHAPTER 40

# Supported XPath syntax

This chapter contains the following topics:

1. [Overview](#)
2. [Example expressions](#)
3. [Syntax specifications for XPath expressions](#)
4. [Java API](#)

## 40.1 Overview

The XPath notation used in TIBCO EBX® must conform to the *abbreviated syntax* of the [XML Path Language \(XPath\) Version 1.0](#) standard, with certain restrictions. This document details the abbreviated syntax that is supported.

## 40.2 Example expressions

The general XPath expression is:

*path[predicate]*

### ***Absolute path***

*/library/books/*

### ***Relative paths***

*./Author*

*../Title*

### ***Root and descendant paths***

*//books*

### ***Table paths with predicates***

*../../books/[author\_id = 0101 and (publisher = 'harmattan')]*

*/library/books/[not(publisher = 'dumesnil')]*

### ***Complex predicates***

*starts-with(col3,'xxx') and ends-with(col3,'yyy') and osd:is-not-null(./col3))*

```
contains(col3 , 'xxx') and ( not(col1=100) and date-greater-than(col2, '2007-12-30') )
```

### ***Predicates with parameters***

`author_id = $param1 and publisher = $param2` where the parameters `$param1` and `$param2` refer respectively to 0101 and 'harmattan'

`col1 < $param1 and col4 = $param2` where the parameters `$param1` and `$param2` refer respectively to 100 and 'true'

`contains(col3,$param1) and date-greater-than(col2,$param2)` where the parameters `$param1` and `$param2` refer respectively to 'xxx' and '2007-12-30'

#### **Note**

The use of this notation is restricted to the Java API since the parameter values can only be set by the method `Request.setXPathParameterAPI` of the Java API.

### ***Predicates on label***

```
osd:label(./delivery_date)='12/30/2014' and ends-with(osd:label(../adress), 'Beijing - China')
```

### ***Predicates on record label***

```
osd:contains-record-label('dumesnil') or osd:contains-record-label('harmattan')
```

### ***Predicates for validation search***

```
osd:has-validation-item()
```

```
osd:has-validation-item('error,info')
```

```
osd:contains-validation-message('xxx')
```

```
osd:contains-validation-message('xxx', 'info,warning')
```

#### **Note**

- XPath functions for validation search cannot be used on XPath predicates defined on associations and foreign key filters.
- The predicates `osd:label`, `osd:contains-record-label` and `osd:contains-validation-message` are localized. The locale can be set by the methods of the Java API `Request.setLocaleAPI` or `Request.setSessionAPI`.

#### **Attention**

To ensure that the search is performed on an up-to-date validation report, it is necessary to perform an explicit validation of the table just before using these predicates.

# 40.3 Syntax specifications for XPath expressions

## Overview

| Expression                    | Format                                                       | Example             |
|-------------------------------|--------------------------------------------------------------|---------------------|
| XPath expression              | <i>&lt;container path&gt;</i> [ <i>predicate</i> ]           | /books[title='xxx'] |
| <i>&lt;container path&gt;</i> | <i>&lt;absolute path&gt;</i> or <i>&lt;relative path&gt;</i> |                     |
| <i>&lt;absolute path&gt;</i>  | /a/b or //b                                                  | //books             |
| <i>&lt;relative path&gt;</i>  | ../b, ./b or b                                               | ../b                |

## Predicate specification

| Expression                     | Format                                                                                                                                                                                                                                                                                                                                                            | Notes/Example                                                                |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| <predicate>                    | Example: A and (B or not(C)) A,B,C:<br><atomic expression>                                                                                                                                                                                                                                                                                                        | Composition of: logical operators parentheses, not() and atomic expressions. |
| <atomic expression>            | <path><comparator><criterion> or<br>method(<path>,<criterion>)                                                                                                                                                                                                                                                                                                    | royalty = 24.5<br>starts-with(title, 'Johnat')<br>booleanValue = true        |
| <path>                         | <relative path> or <i>osd:label</i> (<relative path>)                                                                                                                                                                                                                                                                                                             | Relative to the table that contains it:<br>../authorstitle                   |
| <comparator>                   | <boolean comparator>, <numeric comparator> or <string comparator>                                                                                                                                                                                                                                                                                                 |                                                                              |
| <boolean comparator>           | = or !=                                                                                                                                                                                                                                                                                                                                                           |                                                                              |
| <numeric comparator>           | =, !=, <, >, <=, or >=                                                                                                                                                                                                                                                                                                                                            |                                                                              |
| <string comparator>            | =                                                                                                                                                                                                                                                                                                                                                                 |                                                                              |
| <method>                       | <date method>, <string method>,<br><i>osd:is-null</i> method or <i>osd:is-not-null</i> method                                                                                                                                                                                                                                                                     |                                                                              |
| <date, time & dateTime method> | date-less-than, date-equal or date-greater-than                                                                                                                                                                                                                                                                                                                   |                                                                              |
| <string method>                | matches, starts-with, ends-with, contains, <i>osd:is-empty</i> , <i>osd:is-not-empty</i> , <i>osd:is-empty-or-nil</i> , <i>osd:is-neither-empty-nor-nil</i> , <i>osd:is-equal-case-insensitive</i> , <i>osd:starts-with-case-insensitive</i> , <i>osd:ends-with-case-insensitive</i> , <i>osd:contains-case-insensitive</i> , or <i>osd:contains-record-label</i> |                                                                              |
| <criterion>                    | <boolean criterion>, <numeric criterion>, <string criterion>, <date criterion>, <time criterion>, or <dateTime criterion>                                                                                                                                                                                                                                         |                                                                              |
| <boolean criterion>            | true, false                                                                                                                                                                                                                                                                                                                                                       |                                                                              |
| <numeric criterion>            | An integer or a decimal                                                                                                                                                                                                                                                                                                                                           | -4.6                                                                         |
| <string criterion>             | Quoted character string                                                                                                                                                                                                                                                                                                                                           | 'azerty'                                                                     |

| Expression           | Format                                                                     | Notes/Example         |
|----------------------|----------------------------------------------------------------------------|-----------------------|
| <date criterion>     | Quoted and formatted as 'yyyy-MM-dd'                                       | '2007-12-31'          |
| <time criterion>     | Quoted and formatted as 'HH:mm:ss' or 'HH:mm:ss.SSS'                       | '11:55:00'            |
| <dateTime criterion> | Quoted and formatted as 'yyyy-MM-ddTHH:mm:ss' or 'yyyy-MM-ddTHH:mm:ss.SSS' | '2007-12-31T11:55:00' |

## XPath 1.0 formula

It is possible to use an XPath 1.0 formula in the criterion value part of an atomic predicate expression (right-hand side).

For example, instead of `[./a=3]`, you may use the expression `[./a=(floor(./d)+ 2.0)]`.

Due to the strong dependence of predicates on the data model node and the node type of the criterion, the path portion of the atomic predicate expression (left-hand side) must be a node path and cannot be an XPath formula. For example, the expression `/table[floor(./a) > ceiling(./d)]` is not valid.

## Predicate on label

The `osd:label()` function can be applied to the path portion of the atomic predicate, in order to resolve the predicate on the label instead of the value. In this case, only string operators and string criteria can be used, i.e. `ends-with(osd:label(./price), '99')`.

A predicate on label is localized, so the criterion must be expressed in the same locale as the predicate-filtered request. For example: `request.setLocale(Locale.FRENCH); request.setXPathFilter("osd:label(./delivery_date)='30/12/2014'");`

### Note

It is forbidden to use the `osd:label` function if the right part of the predicate is a contextual value.

### Note

If the `osd:label` function is used in a data model, for example on a selection or in the filter predicate of a table reference node, the default locale of the data model (as defined in its module declaration) must be used for the criterion format (even though this is generally not recommended).

See also `SchemaNode.displayOccurrence`<sup>API</sup>

## Contextual values

For predicates that are relative to a selected node, the criterion value (that is, the right-hand side of the predicate) can be replaced with a contextual path using the syntax `${<relative-path>}` where *<relative-path>* is the location of the element relative to the selected node.

### Note

When calling a method, the criterion is the second parameter, and the first parameter cannot be a relative value.

## Aggregated lists

For predicates on aggregated lists, the predicate returns true regardless of the comparator if one of the list elements verifies the predicate.

### Note

Special attention must be paid to the comparator `!=`. For example, for an aggregated list, `./list != 'a'` is not the same as `not(./list = 'a')`. Where the list contains the elements (e1,e2,...), the first predicate is equivalent to `e1 != 'a' or e2 != 'a' ...`, while the second is equivalent to `e1 != 'a' and e2 != 'a' ...`.

## 'Null' values

Null values must be explicitly treated in a predicate using the operators `osd:is-null` and `osd:is-not-null`.

For example, `/root/products[./price<100]` or `/root/products[./price!=100]` will not return any products whose prices are not set (null). For the latter case to return unset values as well, the predicate must instead be: `/root/products[./price!=100 or osd:is-null(./price)]`.

## How to manage single and double quotes in literal expressions

By default, a literal expression is delimited by single quotes (`'`). If the literal expression contains single quotes and no double quotes, the expression must be delimited by double quotes (`"`). If the literal expression contains both single and double quotes, the single quotes must be doubled.

The method `XPathExpressionHelper.encodeLiteralStringWithDelimitersAPI` in the Java API handles this.

### Examples of using `encodeLiteralStringWithDelimiters`

| Value of Literal Expression            | Result of this method                     |
|----------------------------------------|-------------------------------------------|
| Coeur                                  | 'Coeur'                                   |
| Coeur d'Alene                          | "Coeur d'Alene"                           |
| He said: "They live in Coeur d'Alene". | 'He said: "They live in Coeur d''Alene".' |

## Extraction of foreign keys

In EBX®, the foreign keys are grouped into a single field with the `osd:tableRef` [p 500] declaration.

The standard XPath syntax has been extended so as to extract the value of any targeted primary key field.

### Example

If the table `/root/tableA` has an `osd:tableRef` field named 'fkB' whose target is `/root/tableB` and the primary key of `tableB` has two fields, `id` of type `xs:int` and `date` of type `xs:date`, then the following expressions would be valid:

- `/root/tableA[ fkB = '123|2008-01-21' ]`, where the string "123|2008-01-21" is a representation of the entire primary key value.  
See **Syntax of the internal String representation of primary keys** `PrimaryKey.syntaxAPI` for more information.
- `/root/tableA[ fkB/id = 123 and date-equal(fkB/date, '2008-01-21') ]`, where this predicate is a more efficient equivalent to the one in the previous example.
- `/root/tableA[ fkB/id >= 123 ]`, where any number operator could be used, as the targeted primary key field is of type `xs:int`.
- `/root/tableA[ date-greater-than( ./fkB/date, '2007-01-01' ) ]`, where any date operator could be used, as the targeted primary key field is of type `xs:date`;
- `/root/tableA[ fkB = "" ]` is not valid as the targeted primary key has two columns.
- `/root/tableA[ osd:is-null(fkB) ]` checks if a foreign key is null (not defined).

## 40.4 Java API

Using the XPath in the Java API:

In the Java API, the `xPathFilter` class allows to define XPath predicates and to execute requests on them.

The `xPathExpressionHelper` class provides utilitarian methods to handle XPath predicates and paths.





---

# Localization

---

## CHAPTER 41

# Labeling and localization

This chapter contains the following topics:

1. [Overview](#)
2. [Value formatting policies](#)
3. [Syntax for locales](#)

## 41.1 Overview

TIBCO EBX® offers the ability to handle the labeling and the internationalization of data models.

### ***Localizing user interactions***

In EBX®, language preferences can be set for two scopes:

1. Session: Each user can select a default locale from the user pane.
2. The EBX® main configuration file, named `ebx.properties` by default. See [Extending TIBCO EBX® internationalization](#) [p 237] for more information.

### ***Textual information***

In EBX®, most master data entities can have a label and a description, or can correspond to a user message. For example:

- Dataspaces, snapshots and datasets can have their own label and description. The label is independent of the unique name, so that it remains localizable and modifiable;
- Any node in the data model can have a static label and description;
- Values can have a static label when they are enumerated;
- Validation messages can be customized, and permission restrictions can provide text explaining the reason;
- Each record is dynamically displayed according to its content, as well as the context in which it is being displayed (in a hierarchy, as a foreign key, etc.);

All this textual information can be localized into the locales that are declared in `ebx.properties`.

See also

[Labels and messages](#) [p 533]

[Tables declaration](#) [p 495]

[Foreign keys declaration](#) [p 500]

## 41.2 Value formatting policies

When a value is displayed to the user, it is formatted according to its type and the formatting policy of the current locale. For example, a date will be displayed in some locales as "dd/MM/yyyy" and "MM/dd/yyyy" in others.

A formatting policy is used to define how to display the values of [simple types](#) [p 482].

For each locale declared in `ebx.properties`, its formatting policy is configured in a file located at `/WEB-INF/ebx/{locale}/frontEndFormattingPolicy.xml`. For instance, to define the formatting policy for Greek (el), the engine looks for the following path in the module:

```
/WEB-INF/ebx/el/frontEndFormattingPolicy.xml
```

If the corresponding file does not exist, the formatting policy is looked up in the class-path of EBX®. If the locale-specific formatting policy is not found, the formatting policy of `en_us` is applied.

The content of the file `frontEndFormattingPolicy.xml` is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<formattingPolicy xmlns="urn:ebx-schemas:formattingPolicy_1.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:ebx-schemas:formattingPolicy_1.0 ../schema/ebx-reserved/formattingPolicy_1.0.xsd">
  <date pattern="dd/MM" />
  <time pattern="HH:mm:ss" />
  <dateTime pattern="dd/MM/yyyy HH:mm" />
  <decimal pattern="00,00,00.000" groupingSeparator="|" decimalSeparator="^"/>
  <int pattern="000,000" groupingSeparator=" " />
</formattingPolicy>
```

The elements `date`, `dateTime` and `time` are mandatory.

The group and decimal separators that appear in the formatted numbers can be modified by defining the attributes `groupingSeparator` and `decimalSeparator` for the elements `decimal` and `int`.

## 41.3 Syntax for locales

There are two ways to express a locale:

1. The XML recommendation follows the [IETF BCP 47](#) recommendation, which uses a hyphen '-' as the separator.
2. The Java specification uses an underscore '\_' instead of a hyphen.

In any XML file (XSD, formatting policy file, etc.) read by EBX®, either syntax is allowed.

For a web path, that is, a path within the web application, only the Java syntax is allowed. Thus, formatting policy files must be located in directories whose locale names respect the Java syntax.

See also [Extending TIBCO EBX® internationalization](#) [p 237]



## CHAPTER 42

# Extending TIBCO EBX® internationalization

This chapter contains the following topics:

1. [Overview of the native EBX® localization](#)
2. [Extending EBX® user interface localization](#)
3. [Localized resources resolution](#)
4. [Known limitations](#)

## 42.1 Overview of the native EBX® localization

By default, the EBX® built-in user interface is provided in English (en-US) and French (fr-FR).

Localization consists of a formatting policy and a set of message files (resource bundle):

- For English, localization is provided by a formatting policy and a set of message files with no locale defined,
- For French, localization is provided by a formatting policy and a set of message files with locale set to "fr".

EBX® provides an option to add locales in order to extend the localization of the user interface and to internationalize the documentation of data models and associated services.

## 42.2 Extending EBX® user interface localization

EBX® supports the localization of its user interface into any compatible language and region.

**Note**

Currently, Latin & Cyrillic characters are supported. Locales that use other character sets may be usable, but are not supported.

### ***Adding a new locale***

In order to add a new locale, the following steps must be followed:

- Declare the new locale in the EBX® main configuration file. For example:  
ebx.locales.available=en-US, fr-FR, xx

- The first locale is always considered the default.
- The built-in locales, en-US and fr-FR, can be removed if required.

See [Configuring EBX® localization](#) [p 353].

- Deploy the following files in the EBX® class-path:
  - A formatting policy file, named `com.orchestranetworks.i18n.frontEndFormattingPolicy_xx.xml`,
  - A set of localized message files (\*\_xx.mxml) in a resource bundle.

#### Note

The files must be ending with ".mxml".

## 42.3 Localized resources resolution

Since version 5.7.0, localized resources are resolved on a locale-proximity base, with the following lookup mechanism:

- resourceName + "\_" + language + "\_" + country + "\_" + variant + ".mxml"
- resourceName + "\_" + language + "\_" + country + ".mxml"
- resourceName + "\_" + language + ".mxml"
- resourceName + ".mxml"

#### Note

The resolution is done at the localized message level. It is therefore possible to define one or more files for a locale that only includes messages for which specific localization is required.

## 42.4 Known limitations

### ***Non extendable materials***

Localization of the following cannot be extended:

- EBX® product documentation,
- EBX® HTML editor and viewer.

---

# Persistence

---

## CHAPTER 43

# Overview of persistence

This chapter describes how master data, history, and replicated tables are persisted. A given table can employ any combination of master data persistence mode, historization, and replication.

While all persisted information in TIBCO EBX® is ultimately stored as relational tables in the underlying database, whether it is in a form that is accessible outside of EBX® depends on if it is in mapped mode.

**Note**

The term [mapped mode](#) [p 241] refers to any tables that are stored as-is, and thus whose contents can be accessed directly in the database.

This chapter contains the following topics:

1. [Persistence of managed master data](#)
2. [Historization](#)
3. [Replication](#)
4. [Mapped mode](#)

## 43.1 Persistence of managed master data

Data that is modeled in and governed by the EBX® repository can be persisted in one of two modes, semantic (default) or relational, as specified in its underlying data model. Distinct tables defined in either mode can co-exist and collaborate within the same EBX® repository.

For a comparison between relational mode and semantic mode, see the chapter [Overview of modes](#) [p 243]

## 43.2 Historization

Master data tables can activate historization in order to track modifications to their data, regardless of whether they are persisted in semantic or relational mode, and whether they are replicated.

The history itself is in mapped mode, meaning that it can potentially be consulted directly in the underlying database.

See also [History](#) [p 249]



## 43.3 Replication

Replication enables direct SQL access to tables of master data, by making a copy of data in the repository to relational table replicas in the database. Replication can be enabled on any table regardless of whether it is persisted in semantic or relational mode, and whether it has history activated.

The replica tables are persisted in mapped mode, as their primary purpose is to make master data accessible to direct queries outside of EBX®.

See also [Replication](#) [p 257]

## 43.4 Mapped mode

### ***Overview of mapped mode***

Mapped mode refers to cases where tables are persisted in the underlying relational database in a format that allows their data to be accessed directly, outside of EBX®. Master data modeled in relational mode, history, and replica tables are all examples of tables in mapped mode.

All cases of mapped mode involve automatic alterations of the database schema (the database tables, indexes, etc.) when necessary, by automatically executing required DDL statements in the background. Such procedures are always triggered at data model compilation time and the data model compilation report notifies of any resulting errors.

Another general consideration regarding mapped modes is that, in most cases, when a data model entity is removed, its corresponding database object is not deleted immediately. Instead, it is marked as disabled, which leaves the possibility of later re-enabling the object. In order to definitively drop the object and its associated data and resources from the database, it must be marked for purge. The removal then takes place during the next global purge.

See also

[Database mapping administration](#) [p 407]

[Data model evolutions](#) [p 263]

### ***Data model restrictions due to mapped mode***

Due to the nature of persisting directly in the underlying database, some restrictions apply to all tables stored in mapped mode:

- [Limitations of supported databases](#) [p 311]
- Unlimited-length strings: All string fields, except foreign keys, of type `xs:string`, its derived types, and `xs:anyURI` must define a 'maxLength' or 'length' facet. Since a foreign key field is composed of the final primary key field(s) of its target table(s), this facet requirement applies to each of those final primary key fields instead of the foreign key field itself. Additionally, limitations of the underlying database concerning the maximum length of its character types apply, such as `VARCHAR` and `NVARCHAR2`.
- Large lists of columns might not be indexable. Example for Oracle: the database enforces a limit on the maximum cumulated size of the columns included in an index. For strings, this size also depends on the character set. If the database server fails to create the index, you should consider redesigning your indexes, typically by using a shorter length for the concerned columns, or by

including fewer columns in the index. The reasoning is that an index leading to this situation would have headers so large that it could not be efficient anyway.

- Fields of type `type="osd:password"` are ignored.
- Terminal complex types are supported; however, they cannot be globally set to `null` at record-level.

More generally, tables in mapped mode are subject to any limitations of the underlying RDBMS. For example, the maximum number of columns in a table applies (1000 for Oracle 11g R2, 1600 for PostgreSQL). Note that a history table contains twice as many fields as declared in the schema (one functional field, plus one generated field for the operation code).

Data model evolutions may also be constrained by the underlying RDBMS, depending on the existing data model.

See also [Data model evolutions](#) [p 263]

## CHAPTER 44

# Relational mode

This chapter contains the following topics:

1. [Overview of modes](#)
2. [Enabling relational mode for a data model](#)
3. [Validation](#)
4. [SQL access to data in relational mode](#)
5. [Limitations of relational mode](#)

## 44.1 Overview of modes

### ***Semantic mode explained***

Semantic mode offers all TIBCO EBX® advanced features of master data management, in particular, dataspaces, dataset inheritance, and inherited fields.

Semantic mode is the default mode for persisting the data governed by the EBX® repository. Data models are in semantic mode unless [relational mode](#) [p 244] is explicitly [specified](#) [p 244].

Internally, the master data managed in semantic mode is represented as standard XML, which complies with the XML Schema Document of its data model. The XML representation is additionally compressed and segmented for storage into generic relational database tables. This mode provides efficient data storage and access, including for:

- Dataspaces: no data is duplicated when creating a child dataspace, and
- Inheritance: no data is duplicated when creating an inherited instance.

Semantic mode also makes it possible to maintain an unlimited number of datasets for each data model, organized into an unlimited number of dataspaces or snapshots. This can be done with no impact on the database schema.

As this mode only uses common, generic internal tables, modifications to the structure of the data model also never impact the database schema. Data model evolutions only impact the content of the generic database tables.

See also

[dataspaces](#) [p 88]

[dataset inheritance](#) [p 269]

[inherited fields](#) [p 270]

## Relational mode explained

Relational mode, which is a mapped mode, persists master data directly into the database. The primary function of relational mode is to be able to benefit from the performance and scalability capabilities of the underlying relational database. However, relational mode does not support the advanced governance features offered by semantic mode.

For some cases where the management advantages of semantic mode are not necessary, such as "current time" tables, or tables that are regularly updated by external systems, the performance gains offered by relational mode may be more valuable.

Generally, when a dataset is in relational mode, every table in this dataset has a corresponding table in the database and every field of its data model is mapped to a relational table column.

## Direct comparison of semantic and relational modes

This table summarizes the differences between the two persistence modes:

|                       | Semantic mode                                                              | Relational mode                                                                                       |
|-----------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Dataspaces            | Yes                                                                        | No                                                                                                    |
| Dataset inheritance   | Yes                                                                        | No                                                                                                    |
| Inherited fields      | Yes                                                                        | No                                                                                                    |
| Data model            | All features are supported.                                                | Some restrictions, see <a href="#">Data model restrictions for tables in relational mode</a> [p 248]. |
| Direct SQL reads      | No                                                                         | Yes, see <a href="#">SQL reads</a> [p 247].                                                           |
| Direct SQL writes     | No                                                                         | Yes, but only under precise conditions, see <a href="#">SQL writes</a> [p 247].                       |
| Data validation       | Yes, enables tolerant mode.                                                | Yes, some constraints become blocking, see <a href="#">Validation</a> [p 245].                        |
| Transactions          | See <a href="#">Concurrency and isolation levels</a> [p 468].              | See <a href="#">Concurrency and isolation levels</a> [p 468].                                         |
| Data model evolutions | See <a href="#">Data model evolutions</a> [p 263] in the Reference Manual. | See <a href="#">Data model evolutions</a> [p 263] in the Reference Manual.                            |

## 44.2 Enabling relational mode for a data model

The data model declares that it is in relational mode. Due to the necessary restrictions of relational mode, such as not having child dataspace or snapshots, a specific relational dataspace must be provided, to which the data model will be published. Relational dataspace do not allow creating sub-dataspace or snapshots.

Example of a relational mode declaration:

```
<xs:schema>
  <xs:annotation>
    <xs:appinfo>
      <osd:relationalMode>
        <dataSpace>aDataSpaceKey</dataSpace>
        <dataSet>aDataSetReference</dataSet>
        <tablesPrefix>aPrefixForTablesInRDBMS</tablesPrefix>
      </osd:relationalMode>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:schema>
```

with the elements:

| Element      | Description                                                                                                                                                                                         | Required |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| dataSpace    | Specifies the dataspace where the data model must be published. This dataspace must itself be in relational mode. No dataspace or snapshot can be created from a dataspace declared in such a mode. | Yes      |
| dataSet      | Specifies the dataset where the data model must be published.                                                                                                                                       | Yes      |
| tablesPrefix | Specifies the common prefix used for naming the generated tables in the database.                                                                                                                   | Yes      |

## 44.3 Validation

This section details the impact of relational mode on data validation.

### **Structural constraints**

Some EBX® data model constraints will generate a "structural constraint" on the underlying RDBMS schema for relational mode and also if [table history is activated](#) [p 249]. This concerns the following facets:

- facets `xs:maxLength` and `xs:length` on string elements;
- facets `xs:totalDigits` and `xs:fractionDigits` on `xs:decimal` elements.

Databases do not support as tolerant a validation mode as EBX®. Hence, the above constraints become *blocking constraints*. A blocking constraint means that updates are rejected if they do not comply. Additionally, such constraints are no longer checked during validation process, except for foreign key constraints under some circumstances (see [Foreign key blocking mode](#) [p 245]). When a transaction does not comply with a blocking constraint, it is cancelled and a `ConstraintViolationException`<sup>API</sup> is thrown.

See also [Blocking and non-blocking constraints](#) [p 524]

### **Foreign key blocking mode**

In order to reduce validation time, foreign key constraints are automatically set in blocking mode if:

- The foreign key constraints are defined on a table in relational mode,
- The foreign key constraints are defined on a table in semantic or relational mode referencing a table in relational mode.

For these constraints, blocking mode implies that attempting the following actions will result in a `ConstraintViolationException`<sup>API</sup>:

- Deleting a record referenced by a foreign key constraint,
- Deleting an instance referenced by a foreign key constraint,
- Closing a dataspace referenced by a foreign key constraint.

However, it is possible to overwrite this behavior by setting a specific control policy. See [Blocking and non-blocking constraints](#) [p 524] for more information.

In order to ensure the integrity of foreign key constraints after direct SQL writes that bypass the EBX® governance framework, the foreign key constraints will be validated on the following cases:

- On the first explicit validation through the user interface or API,
- On the first explicit validation through the user interface or API after refreshing the schema,
- On the first explicit validation through the user interface or API after resetting the validation report of a dataset in the user interface.

**Important:**

- Blocking aspect of the foreign key constraint does not concern filters that may be defined. That is, a foreign key constraint is non-blocking if a referenced record exists but does not satisfy a foreign key filter. In this case, updates are not rejected and thus an error will be added to the validation report.
- Foreign key constraints are not in blocking mode upon archive import. Indeed, all blocking constraints, excepted structural constraints, are always disabled when importing archives. This allows flexibility upon archive import where under certain circumstances the import of foreign keys referencing records that are not yet imported must be tolerant.

### ***Constraints on the whole table***

Programmatic **constraints** `Constraint`<sup>API</sup> are checked on each record of the table at validation time. If the table defines millions of records, this becomes a performance issue. It is then recommended to define a **table-level constraint** `ConstraintOnTable`<sup>API</sup>.

In the case where it is not possible to define such a table-level constraint, it is recommended to at least define a **local or explicit dependency** `DependenciesDefinitionContext.dependencies`<sup>API</sup>, so as to reduce the cost of incremental validation.

See also `ConstraintOnTable`<sup>API</sup>

## **44.4 SQL access to data in relational mode**

This section describes how to directly access the data in relational mode, through SQL.

See also [SQL access to history](#) [p 252]

## Finding the table in the database

For every EBX® table in relational mode, a corresponding table is generated in the RDBMS. Using the EBX® user interface, you can find the name of this database table by clicking on the documentation pane of the table. The table in the RDBMS contains the following additional technical columns:

- T\_CREATOR\_ID, reference of the user that has created the record
- T\_LAST\_USER\_ID, reference of the last user that has updated the record
- T\_CREATION\_DATE, creation date of the record
- T\_LAST\_WRITE, time of last modification of the record

Columns T\_CREATOR\_ID and T\_LAST\_USER\_ID have a maximum length of 320 characters. As a consequence, identifiers of users defined in EBX® directory must have a maximum length less than this limit.

## SQL reads

Direct SQL reads are possible in well-managed, preferably short-lived transactions. However, for such accesses, EBX® permissions are not taken into account. As a result, applications given allowed to perform reads must be trusted through other authentication processes and permissions.

## SQL writes

Direct SQL writes bypass the governance framework of EBX®. Therefore, they must be used with **extreme caution**. They could cause the following situations:

- failure to historize EBX® tables;
- failure to execute EBX® triggers;
- failure to verify EBX® permissions and constraints;
- modifications missed by the incremental validation process;
- losing visibility on EBX® semantic tables, which might be referenced by foreign keys.

Consequently, direct SQL writes are to be performed *if, and only if, all the following conditions are verified*:

- The written tables are not historized and have no EBX® triggers.
- The application performing the writes can be fully trusted with the associated permissions, to ensure the integrity of data. Specifically, the integrity of foreign keys (osd:tableRef) must be preserved at all times. See [Foreign key blocking mode](#) [p 245] for more information.
- The application server running EBX® is shut down *whenever writes are performed*. This is to ensure that incremental validation does not become out-of-date, which would typically occur in a batch context.

## 44.5 Limitations of relational mode

The relational mode feature is fully functional, but has some known limitations, which are listed below. If using relational mode, it is strongly recommended to read these limitations carefully and to contact the TIBCO EBX® Support team at <https://support.tibco.com> in case of questions.

See [Supported databases](#) [p 311] for the databases on which relational mode is supported.

## ***Data model restrictions for tables in relational mode***

Some restrictions apply to data models in relational mode:

- [Data model restrictions due to mapped mode](#) [p 241]
- [Aggregated lists](#) [p 492] are not supported in relational tables. Such a schema will cause a compilation error.
- User-defined attributes on relational tables result in data model compilation errors.
- [Dataset inheritance](#) [p 269].
- [Inherited fields](#) [p 270].
- Programmatic constraints, since the computation cost of validation would be too high. However, **constraints on tables** `ConstraintOnTableAPI` remain available.

Schema evolutions may also be constrained by the underlying RDBMS, depending on the data already contained in the concerned tables.

See also [Data model evolutions](#) [p 263]

## ***Other limitations of relational mode***

- [Limitations of mapped mode](#) [p 241]
- From a dataspace containing datasets in relational mode, it is not possible to create child dataspace and snapshots.
- For D3, it is not possible to broadcast a dataspace defined in relational mode.
- For very large volumes of data, the validation will show poor performance if the relational table declares any of these features: `osd:function`, `osd:select`, `osd:uiFilter`, `osd:tableRef/filter`. Additionally, a sort cannot be applied on a `osd:function` column.
- It is not possible to set the `AdaptationValue.INHERIT_VALUE` to a node belonging to a data model in relational mode.



# CHAPTER 45

---

## History

This chapter contains the following topics:

1. [Overview](#)
2. [Configuring history](#)
3. [History views and permissions](#)
4. [SQL access to history](#)
5. [Impacts and limitations of historized mode](#)

### 45.1 Overview

History is a feature allowing to track all data modifications on a table (records creation, update and deletion).

It is an improvement on the [XML audit trail](#) [p 421]. XML audit trail is still activated by default; it can be safely deactivated if history is enabled for the relevant tables.

#### See also

[History](#) [p 26]

[Relational mode](#) [p 243]

[Replication](#) [p 257]

[Data model evolutions](#) [p 263]

### 45.2 Configuring history

In order to activate historization for a table, a history profile has to be set for the table in the data model. This section describes history profiles and the way they are associated with tables.

#### ***Configuring history in the repository***

A history profile specifies when the historization is to be created. In order to edit history profiles, select *Administration > History and logs*.

A history profile is identified by a name and defines the following information:

- An internationalized label.

- A list of dataspace (branches) for which history is activated. It is possible to specify whether direct children and/or all descendants should also be concerned.

Some profiles are already created when installing the repository. These profiles can neither be deleted nor modified.

| Profile Id          | Description                                                                                                                                                           |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ebx-referenceBranch | This profile is activated only on the reference dataspace.                                                                                                            |
| ebx-allBranches     | This profile is activated on all dataspaces.                                                                                                                          |
| ebx-instanceHeaders | This profile historizes dataset headers. However, this profile will only be setup in a future version, given that the internal data model only defines dataset nodes. |

## Configuring history in the data model

### Activating table history

History can be activated on a table either through the data model assistant, or by editing the underlying data model.

To activate history by editing the data model, a history profile should be declared on the table using the `historyProfile` element.

```
<osd:table>
  <primaryKeys>/key</primaryKeys>
  <historyProfile>historyProfileForProducts</historyProfile>
</osd:table>
```

The data model assistant allows you to view the historization profiles defined in the repository.

Historization must be activated for each table separately. See [model design](#) [p 478] documentation for more details.

### Disabling history on a specific field or group

For a historized table, the default behavior is to historize all its supported elements (see [Impacts and limitations of historized mode](#) [p 254]).

It is possible to disable history for a given field or group, either through the data model assistant, or by editing the underlying data model.

To disable the history of a field or group by editing the data model, use the element `osd:history` with the attribute `disable="true"`.

```
<xs:element name="longDescription" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:history disable="true" />
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

To disable the history of a field or group through the data model assistant, use the `History` property in the Advanced properties of the element.

When this property is defined on a group, history is disabled recursively for all its descendants. Once a group disables history, it is not possible to specifically re-enable history on a descendant.

#### Note

If the table containing the field or group is not historized, this property will not have any effect.  
It is not possible to disable history for primary key fields.

## Integrity

If problems are detected at data model compilation, warning messages or error messages will be added to the validation report associated with this data model. Furthermore, if any error is detected, each associated instance (dataset) will be inaccessible. The most common error cases are the following:

- A table references a profile that is not defined in the repository.
- A history profile that is referenced in the data model mentions a non-defined or closed dataspace in the current repository.

#### Note

Deploying a data model on a repository that does not have the expected profiles requires the administrator to add them.

## 45.3 History views and permissions

### Table history view

When the history has been activated on a table in the data model, it is possible to access the history view from various locations in the user interface: record, selection of records, table and dataset.

The next section explains how permissions are resolved.

For more information, see [table history view](#) [p 27] section. To access the table history view from Java, the method `AdaptationTable.getHistoryAPI` must be invoked.

### Permissions for table history

Data permissions are also applied to data history. History permissions are resolved automatically as the most restricted permission between data permissions and *read-only* access right.

This is true for user-defined permission rules and also for programmatic permission rules.

When defining a programmatic rule, it may be required to distinguish between the functional dataset context and the history view context, either because the expected permissions are not the same, or because some fields are not present in the history structure. This is the case for dataset fields, computed values and [fields for which history has been disabled](#) [p 250]. The methods `Adaptation.isHistoryAPI` and `AdaptationTable.getHistoryAPI` can then be used in the programmatic rule in order to implement specific behavior for history.

### Transaction history views

The transaction history view gives access to the executed transactions, independently of a table, a dataset or a data model, directly from the user interface.

To see the 'Transaction history' table, navigate to the Administration area and select 'History and logs' using the down arrow menu in the navigation pane. Transaction history can also be accessed from the Dataspaces area by selecting a historized datasource and using the **Actions** menu in the workspace.

For more information, see [transaction history view](#) [p 27].

# 45.4 SQL access to history

This section describes how to directly access the history data by means of SQL.

See also [SQL access to data in relational mode](#) [p 246]

## Access restrictions

The database tables must be accessed only in read-only mode. It is up to the database administrator to forbid write access except for the database user used by TIBCO EBX®, as specified in the section [Rules for the database access and user privileges](#) [p 375].

## Relational schema overview

Here is a description of the history tables in the database.

The database schema contains (see also the diagram in the next section):

|                                  |                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Common and generic tables</b> | <p>The main table is HV_TX; each record of this table represents a transaction. Only transactions that involve at least one historized table are recorded.</p> <p>These common tables are all prefixed by "HV".</p>                                                                                                                                         |
| <b>Specific generated tables</b> | <p>For each historized table, a specific history table is generated. This table contains the history of the data modifications on the table.</p> <p>In the EBX® user interface, the name of this table in database can be obtained by clicking on the table documentation pane (advanced mode). All the specific history tables are prefixed with "HG".</p> |


## Example of a generated history table

In the following example, we are historizing a table called product. Let us assume this table declares three fields in EBX® data model:

Product

- productId: int
- price: int
- beginDate: Date

The diagram below shows the resulting relational schema:



Activating history on this table generates the HG\_product table shown in the history schema structure above. Here is the description of its different fields:

- tx\_id: transaction ID.
- instance: instance ID.
- op: operation type - C (create), U (update) or D (delete).
- productId: productId field value.
- OproductId: operation field for productId, see next section.
- price: price field value.
- Oprice: operation field for price, see next section.
- beginDate: date field value.
- ObeginDate: operation field for beginDate, see next section.

## Combination of operations

If several operations are combined in the same transaction, the operation field is resolved as follows:

- $C + U \rightarrow C$   
 $D + U \rightarrow D$   
 $D + C \rightarrow U$   
 $C + D \rightarrow \{\}$  (no entry in history)

### Values for operation fields

For each functional field, an additional operation field is defined, composed of the field name prefixed by the character o. This field specifies whether the functional field has been modified. It is set to one of the following values:

- `null`: if the functional field value has not been modified (and its value is not INHERIT).
- `M`: if the functional field value has been modified (not to INHERIT).
- `D`: if record has been deleted.

If [inheritance](#) [p 268] is enabled, the operation field can have three additional values:

- `T`: if the functional field value has not been modified and its value is INHERIT.
- `I`: if the functional field value has been set to INHERIT.
- `O`: if the record has been set to OCCULTING mode.

## 45.5 Impacts and limitations of historized mode

The history feature has some impacts and known limitations, which are listed in this section. If using historized mode, it is strongly recommended to read these limitations carefully and to contact Cloud Software Group, Inc. support in case of questions.

### **Validation**

Some EBX® data model constraints become blocking constraints when table history is activated. For more information, see the section [Structural constraints](#) [p 245].

### **Data model restrictions for historized tables**

Some restrictions apply to data models containing historized tables:

- [Data model restrictions due to mapped mode](#) [p 241]
- Limitations exist for two types of aggregated lists: aggregated lists under another aggregated list, and aggregated lists under a terminal group. Data models that contain such aggregated lists can be used, however these lists will be ignored (not historized).
- Computed values are ignored.
- User-defined attributes on historized tables result in data model compilation errors.

Data model evolutions may also be constrained by the underlying RDBMS, depending on the data already contained in the concerned tables.

See also [Data model evolutions](#) [p 263]

### **Other limitations of historized mode**

- No data copy is performed when a table with existing data is activated for history.
- Global operations on datasets are not historized (create an instance and remove an instance), even if they declare a historized table.
- Default labels referencing a non-historized field are not supported for historized tables.

As a consequence, default labels referencing a computed field are not supported for historized tables.

The workaround is to implement the `UILabelRenderer` interface and adapt the label computation for history.

- D3: the history can be enabled in the delivery dataspace of a primary node, but in the delivery dataspace of the replica nodes, the historization features are always disabled.

- Recorded user in history: for some specific operations, the user who performs the last operation and the one recorded in the corresponding history record may be different.

This is due to the fact that these operations are actually a report of the data status at a previous state:

- Archive import: when importing an archive on a dataspace, the time and user of the last operation performed in the child dataspace are preserved, while the user recorded in history is the user who performs the import.
- Programmatic merge: when performing a programmatic merge on a dataspace, the time and user of the last operation performed in the child dataspace are preserved, while the user recorded in history is the user who performs the merge.
- D3: for distributed data delivery feature, when a broadcast is performed, the data from the primary node is reported on the replica node and the time and user of the last operation performed in the child dataspace are preserved, while the user recorded in history is 'ebx-systemUser' who performs the report on the replica node upon the broadcast.





# CHAPTER 46

---

## Replication

This chapter contains the following topics:

1. [Overview](#)
2. [Configuring replication](#)
3. [Accessing a replica table using SQL](#)
4. [Requesting an 'onDemand' replication refresh](#)
5. [Impact and limitations of replication](#)

### 46.1 Overview

Data stored in the TIBCO EBX® repository can be mirrored to dedicated relational tables to enable direct access to the data by SQL requests and views.

Like history and relational mode, this data replication is transparent to end-users and client applications. Certain actions trigger automatic changes to the replica in the database:

- Activating replication at the model-level updates the database schema by automatically executing the necessary DDL statements.
- Data model evolutions that impact replicated tables, such as creating a new column, also automatically update the database schema using DDL statements.
- When using the 'onCommit' refresh mode: updating data in the EBX® repository triggers the associated inserts, updates, and deletions on the replica database tables.

#### See also

[Relational mode](#) [p 243]

[History](#) [p 249]

[Data model evolutions](#) [p 263]

[Repository administration](#) [p 374]

#### Note

*replicated table*: refers to a primary data table that has been replicated

*replica table* (or *replica*): refers to a database table that is the target of the replication

## 46.2 Configuring replication

### *Enabling replication*

To define a *replication unit* on a data model, use the element `osd:replication` under the elements `annotation/appinfo`. Each replication unit specifies tables in a single dataset in a specific dataspace.

The nested elements are as follows:

| Element                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Required |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| <code>name</code>                         | Name of the replication unit. This name identifies a replication unit in the current data model. It must be unique.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Yes      |
| <code>dataSpace</code>                    | Specifies the dataspace relevant to this replication unit. It cannot be a snapshot or a relational dataspace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Yes      |
| <code>dataSet</code>                      | Specifies the dataset relevant to this replication unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Yes      |
| <code>refresh</code>                      | Specifies the data synchronization policy. The possible policies are: <ul style="list-style-type: none"> <li><code>onCommit</code>: The replica table content in the database is always up to date with respect to its source table. Every transaction that updates the EBX® source table triggers the corresponding insert, update, and delete statements on the replica table.</li> <li><code>onDemand</code>: The replication of specified tables is only done when an explicit refresh operation is performed. See <a href="#">Requesting an 'onDemand' replication refresh</a> [p 260].</li> </ul> | Yes      |
| <code>table/path</code>                   | Specifies the path of the table in the current data model that is to be replicated to the database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Yes      |
| <code>table/nameInDatabase</code>         | Specifies the name of the table in the database to which the data will be replicated. This name must be unique amongst all replications units.                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Yes      |
| <code>table/element/path</code>           | Specifies the path of the aggregated list in the table that is to be replicated to the database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Yes      |
| <code>table/element/nameInDatabase</code> | Specifies the name of the table in the database to which the data of the aggregated list will be replicated. This name must be unique amongst all replications units.                                                                                                                                                                                                                                                                                                                                                                                                                                   | Yes      |

For example:

```
<xs:schema>
  <xs:annotation>
    <xs:appinfo>
      <osd:replication>
        <name>ProductRef</name>
        <dataSpace>ProductReference</dataSpace>
        <dataSet>productCatalog</dataSet>
        <refresh>onCommit</refresh>
        <table>
          <path>/root/domain1/tableA</path>
          <nameInDatabase>PRODUCT_REF_A</nameInDatabase>
        </table>
      </osd:replication>
    </xs:appinfo>
  </xs:annotation>
</xs:schema>
```

```

<table>
  <path>/root/domain1/tableB</path>
  <nameInDatabase>PRODUCT_REF_B</nameInDatabase>
  <element>
    <path>/retailers</path>
    <nameInDatabase>PRODUCT_REF_B_RETAILERS</nameInDatabase>
  </element>
</table>
</osd:replication>
</xs:appinfo>
</xs:annotation>
...
</xs:schema>

```

**Notes:**

- See [Data model restrictions for replicated tables](#) [p 261]
- If, at data model compilation, the specified dataset and/or dataspace does not exist in the current repository, a warning is reported, but the replica table is created in the database. Once the specified dataspace and dataset are created, the replication becomes active.
- At data model compilation, if a table replication is removed, or if some of the above properties has changed, the replica table is dropped from the database, and then recreated with the new definition if needed.

***Disabling replication on a specific field or group***

For a replicated table, the default behavior is to replicate all its supported elements (see [Data model restrictions for replicated tables](#) [p 261]).

It is possible to disable replication for a specific field or group, either through the data model assistant, or by editing the underlying data model.

To disable the replication of a field or group by editing the data model, use the element `osd:replication` with the attribute `disable="true"`.

```

<xs:element name="longDescription" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:replication disable="true" />
    </xs:appinfo>
  </xs:annotation>
</xs:element>

```

To disable the replication of a field or group through the data model assistant, use the Replication property in the Advanced properties of the element.

When this property is defined on a group, replication is disabled recursively for all its descendents. Once a group disables replication, it is not possible to specifically re-enable replication on a descendant.

**Note**

If the table containing the field or group is not replicated, this property will not have any effect. It is not possible to disable replication for primary key fields.

## 46.3 Accessing a replica table using SQL

This section describes how to directly access a replica table using SQL.

See also [SQL access to history](#) [p 252]

## ***Finding the replica table in the database***

For every replicated EBX® table, a corresponding table is generated in the RDBMS. Using the EBX® user interface, you can find the name of this database table by clicking on the documentation pane of the table. The table in the RDBMS contains the following additional technical columns:

- T\_CREATOR\_ID, reference of the user that has created the record
- T\_LAST\_USER\_ID, reference of the last user that has updated the record
- T\_CREATION\_DATE, creation date of the record
- T\_LAST\_WRITE, time of last modification of the record

Columns T\_CREATOR\_ID and T\_LAST\_USER\_ID have a maximum length of 320 characters. As a consequence, identifiers of users defined in EBX® directory must have a maximum length less than this limit.

## ***Access restrictions***

The replica database tables must only be directly accessed in read-only mode. It is the responsibility of the database administrator to block write-access to all database users except the one that EBX® uses.

See also [Rules for the database access and user privileges](#) [p 375]

## ***SQL reads***

Direct SQL reads are possible in well-managed, preferably short-lived transactions. However, for such accesses, EBX® permissions are not taken into account. As a result, applications given the privilege to perform reads must be trusted through other authentication processes and permissions.

# **46.4 Requesting an 'onDemand' replication refresh**

The 'onDemand' refresh policy requires an explicit request to refresh the replicated table data.

There are several ways to request a replication refresh:

- **User interface:** In the dataset actions menu, use the action 'Refresh replicas' under the group 'Replication' to launch the replication refresh wizard.
- **Data services:** Use the replication refresh data services operation. See [Replication refresh](#) [p 644] for data services for more information.
- **Java API:** Call the `ReplicationUnit.performRefreshAPT` methods in the `ReplicationUnit` API to launch a refresh of the replication unit.

# **46.5 Impact and limitations of replication**

The replication feature has some known limitations and side-effects, which are listed below. If using replication, it is strongly recommended to read this section carefully and to contact Cloud Software Group, Inc. support in case of questions.

See [Supported databases](#) [p 311] for the databases for which replication is supported.

## Validation

Some EBX® data model constraints become blocking constraints when replication is enabled. For more information, see [Structural constraints](#) [p 245].

## Data model restrictions for replicated tables

Some restrictions apply to data models containing tables that are replicated:

- [Data model restrictions due to mapped mode](#) [p 241]
- Dataset inheritance is not supported for the 'onCommit' refresh policy if the specified dataset is not a root dataset or has not yet been created. See [dataset inheritance](#) [p 269] for more information.
- Field inheritance is also only supported for the 'onDemand' refresh policy. This means that, at data model compilation, an error is reported if the refresh mode is 'onCommit' and the table to be replicated has an inherited field. See [inherited fields](#) [p 270] for more information.
- Computed values are ignored.
- Limitations exist for two types of aggregated lists: aggregated lists under another aggregated list, and aggregated lists under a terminal group. Data models that contain such aggregated lists can be used, however these lists will be ignored (not replicated).
- User-defined attributes are not supported. A compilation error is raised if they are included in a replication unit.

Data model evolutions may also be constrained by the underlying RDBMS, depending on the data already contained in the concerned tables.

See also [Data model evolutions](#) [p 263]

## Database configuration

The refresh operation is optimized to transmit only the rows of the source table that have been modified (with respect to creation and deletion) since the last refresh. However, depending on the volume of data exchanged, this can be an intensive operation, requiring large transactions. In particular, the first refresh operation can concern a large number of rows. It is necessary for the database to be configured properly to allow such transactions to run under optimal conditions.

For instance, with Oracle:

- It is mandatory for the bulk of all replica tables in a replication unit to fit into the 'UNDO' tablespace.
- It is recommended to provide enough space in the buffer cache to allow those transactions to run with minimal disk access.
- It is recommended to provision 'REDO' log groups big enough to avoid those transactions to wait for the 'db\_writer' process.

## Distributed data delivery (D3)

Replication is available on both D3 primary and replica delivery dataspace. On the primary dataspace, the replication behavior is the same as in a standard semantic dataspace, but on replica dataspace, the replicated content is that of the last broadcast snapshot.

In a replica delivery dataspace, some restrictions occur:

- The refresh policy defined in the data model has no influence on the behavior described above: replication always happens on snapshot.
- The action item Refresh replicas is not available.
- It is not allowed to invoke the `ReplicationUnit.performRefreshAPI` method.

See also [D3 overview](#) [p 426]

### ***Other limitations of replication***

- [Limitations of supported databases](#) [p 311]
- For inheritance, a replica record field cannot hold the "inherit value" flag (`AdaptationValue.INHERIT_VALUE`). It only holds the inherited value in such cases. More generally, it is not possible to distinguish inheriting state from overwriting state.

## CHAPTER 47

# Data model evolutions

This chapter describes the modifications that are possible on data models, as well as potential limitations. The restrictions and/or potential impacts of data model evolutions depend on the persistence mode. The principles for each mode are the following:

- Semantic mode: flexible and non-blocking. Can lead to a loss of data; for instance, a primary key definition can freely evolve, but all existing records in any dataspace and snapshot that violate the primary key constraint will no longer be loaded.
- Any mapped mode: restrictive and thus blocking if data exists and if the evolution would violate their integrity according to the new data model.

**Attention**

Whenever the data modeler performs an evolution on the data model, it is important to anticipate the fact that it could lead to a loss of data. In such cases, if existing data must be preserved in some ways, a data migration plan must be set up and operated before the new data model is published or deployed. It can also be noted that data is not destroyed immediately after the data model evolution; in semantic mode, as long as no update is performed on a table whose definition has evolved, if the data model is rolled back to its previous state, then the previous data is retrieved.

**Note**

Certain types of data model evolutions cannot be performed directly in the user interface, and thus the data model must be exported, modified in XSD format, then re-imported. For changes to a data model that impact its configuration, not just its structure, the XSD must be imported into TIBCO EBX® from a module. Otherwise, the configuration modifications are not taken into account.

See also [Mapped mode](#) [p 241]

This chapter contains the following topics:

1. [Types of permitted evolutions](#)
2. [Limitations/restrictions](#)

## 47.1 Types of permitted evolutions

This section describes the possible modifications to data models after their creation.

## ***Model-level evolutions***

The following modifications can be made to existing data models:

- A data model in semantic mode can be declared to be in relational mode. Data should be manually migrated, by exporting then re-importing an XML or archive file.
- Relational mode can be disabled on the data model. Data should be manually migrated, by exporting then re-importing an XML or archive file.
- Replication units can be added to the data model. If their refresh policy is 'onCommit', the corresponding replica tables will be created and refreshed on next schema compilation.
- Replication units can be removed from the data model. The corresponding replica tables will be dropped immediately.
- The data model can be deleted. If it declares replication units, the corresponding replica tables will be dropped immediately. If it is relational or contains historized tables, this change marks the associated mapped tables as disabled. See [Database mapping](#) [p 407] for the actual removal of associated database objects.

## ***Table-level evolutions***

The following modifications can be made to a data model at the table-level:

- A new table can be added. Upon creation, the table can also declare one or more mapped modes.
- An existing table can be deleted. If it declares replication units, the corresponding replica tables will be dropped immediately. If it historized or relational, this change marks the mapped table as disabled. See [Database mapping](#) [p 407] for the actual removal of associated database objects.
- An existing table in semantic mode can be declared to be in relational mode. Data should be manually migrated, by exporting then re-importing an XML or archive file.
- History can be enabled or disabled on a table. History will not take into account the operations performed while it is disabled.
- A table can be renamed. Data should be manually migrated, by exporting then re-importing an XML or archive file, because this change is considered to be a combination of deletion and creation.

## ***Field-level evolutions***

The following modifications can be made to a data model at the field-level:

- A new field can be added.
- An existing field can be deleted. In semantic mode, the data of the deleted field will be removed from each record upon its next update. For a replica table, the corresponding column is automatically removed. In history or relational mode, the field is marked as disabled.
- A field can be specifically disabled from the history or replication which applies to its containing table, by using the attribute `disable="true"`. For a replica table, the corresponding column is automatically removed. For a history table, the column remains but is marked as disabled. See [Disabling history on a specific field or group](#) [p 250] and [Disabling replication on a specific field or group](#) [p 259].
- The facets of a field can be modified, except for the facets listed under [Limitations/restrictions](#) [p 265].



The above-mentioned changes are accepted, but they can lead to a loss of data. Data should be migrated manually, by exporting then re-importing an XML or archive file, since these changes are considered to be a combination of deletion and creation.

- A field can be renamed.
- The type of a field can be changed.

### ***Index-level evolutions***

- An index can be added or renamed.
- An index can be modified, by changing or reordering its fields. In mapped mode, the existing index is deleted and a new one is created.
- An index can be deleted. In mapped mode, a deleted index is also deleted from the database.

## **47.2 Limitations/restrictions**

### **Note**

All limitations listed in this section that affect mapped mode can be worked around by purging the mapped table database resources. For the procedure to purge mapped table database resources, see [Database mapping](#) [p 407].

### ***Limitations related to primary key evolutions***

When a primary key definition is modified:

- In semantic mode, the existing records are only loaded into the cache if they:
  - Respect the uniqueness constraint of the primary key,
  - Comply with the new structure of the primary key.
- In mapped mode, the underlying RDBMS only accepts a primary key modification if all table records respect its uniqueness and non-nullity constraints. In particular, if a table already has existing records:
  - Adding a new field to the primary key requires assigning a default value to this field. Workaround for replicated or relational tables: first add the field, value it for the existing records, then add the field to the primary key.
  - Removing an existing field from the primary key will be rejected if it would cause the existing records to no longer have a unique primary key (assigning a default value makes no change in this case).
  - It is generally not possible to rename a field of the primary key; more formally, it is only possible if the field was not needed for making all primary keys unique. Indeed, renaming a field translates to a combination of deletion and creation; consequently, the operation will be rejected if it would cause the existing records to no longer have a unique primary key (assigning a default value makes no change in this case).

### ***Limitations related to foreign key evolutions***

- When the declaration of a facet `osd:tableRef` is added or modified, or the primary key of the target table of a facet `osd:tableRef` is modified:

- In semantic mode, the existing values for this field are only loaded into the cache if they comply with the new structure of the target primary key.
- In mapped mode, the structure of a foreign key field is set to match that of the target primary key. A single field declaring an `osd:tableRef` constraint may then be split into a number of columns, whose number and types correspond to that of the target primary key. Hence, the following cases of evolutions will have an impact on the structure of the mapped table:
  - declaring a new `osd:tableRef` constraint on a table field;
  - removing an existing `osd:tableRef` constraint on a table field;
  - adding (resp. removing) a column to (resp. from) a primary key referenced by an existing `osd:tableRef` constraint;
  - modifying the type or path for any column of a primary key referenced by an existing `osd:tableRef` constraint.

These cases of evolution will translate to a combination of field deletions and/or creations. Consequently, the existing data should be migrated manually.

### ***Limitations related to field-level evolutions***

- In mapped mode, when a `maxLength`, `length`, `totalDigits` or `fractionDigits` facet is modified: Whether or not this modification is accepted depends on the underlying DBMS, as well as the field type and the contents of the table.  
For example, Oracle will accept changing a `VARCHAR(20)` to a `VARCHAR(50)`, but will only change a `VARCHAR(50)` to a `VARCHAR(20)` if the table does not contain any values over 20 characters long.  
PostgreSQL has the same limitations, but additionally, any modification of a field type (including modifications of its length) will invalidate all related prepared statements, and require restarting the application server.
- When a cardinality of an element is modified:
  - In semantic mode, this change is supported. However, two cases are distinguished:
    - When changing a single element to an aggregated list, the previous single value is preserved and added to the new aggregated list.
    - When changing an aggregated list to a single element, only the last value of the aggregated list is preserved in the single element. Other values are lost.
  - In relational mode, aggregated lists are not supported. An error message is added to the compilation report of the data model if an element is changed to an aggregated list.
  - In historized mode, when changing a single element to an aggregated list, the modification is taken into account, but the previous single value is lost.

---

# Other

---

## CHAPTER 48

# Inheritance and value resolution

This chapter contains the following topics:

1. [Overview](#)
2. [Dataset inheritance](#)
3. [Inherited fields](#)
4. [Optimize & Refactor service](#)

## 48.1 Overview

The principle of inheritance is to mutualize resources that are shared by multiple contexts or entities. TIBCO EBX® offers mechanisms for defining, factorizing and resolving data values: *dataset inheritance* and *inherited fields*.

Furthermore, *functions* can be defined to compute values.

**Note**

Inheritance mechanisms described in this chapter should not be confused with "structural inheritance", which usually applies to models and is proposed in UML class diagrams for example.

See also [Inheritance \(glossary\)](#) [p 25]

### ***Dataset inheritance***

Dataset inheritance is particularly useful when data applies to global enterprise contexts, such as subsidiaries or business partners.

Based on a hierarchy of datasets, it is possible to factorize common data into the root or intermediate datasets and define specialized data in specific contexts.

The dataset inheritance mechanisms are detailed below in [Dataset inheritance](#) [p 269].

### ***Inherited fields***

Contrary to dataset inheritance, which exploits global built-in relationships between datasets, inherited fields exploit finer-grained dependencies that are specific to the data structure. It allows factorizing and specializing data at the business entities-level.

For example, if the model specifies that a 'Product' is associated with a 'FamilyOfProducts', it is possible that some attributes of 'Product' inherit their values from the attributes defined in its associated 'FamilyOfProducts'.

#### Note

When using both inheritance in the same dataset, field inheritance has priority over the dataset one.

## Computed values (functions)

In the data model, it is also possible to specify that a node holds a *computed value*. In this case, the specified JavaBean function will be executed each time the value is requested.

The function is able to take into account the current context, such as the values of the current record or computations based on another table, and to send requests to third-party systems.

See also [Computed values](#) [p 529]

## 48.2 Dataset inheritance

### Dataset inheritance declaration

The dataset inheritance mechanism is declared as follows in a data model:

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:ebxnd="urn:ebx-schemas:binding_1.0">
  <xs:annotation>
    <xs:appinfo>
      <osd:inheritance>
        <dataSetInheritance>all</dataSetInheritance>
      </osd:inheritance>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:schema>
```

The element `osd:inheritance` defines the property `dataSetInheritance` to specify the use of inheritance on datasets based on this data model. The following values can be specified:

- `all`, indicates that inheritance is enabled for all datasets based on the data model.
- `none`, indicates that inheritance is disabled for all datasets based on the data model.

If not specified, the inheritance mechanism is disabled.

### Value lookup mechanism

The dataset inheritance lookup mechanism for values proceeds as follows:

1. If the value is locally defined, it is returned.  
It can be explicitly `null`.
2. Otherwise, looks up the first locally defined value according to the built-in child-to-parent relationship of the dataset in the hierarchy of datasets.
3. If no locally defined value is found, the default value is returned.  
If no default value is defined, `null` is returned.

**Note:** Default values cannot be defined on:

- A single primary key node

- Auto-incremented nodes
- Nodes defining a computed value

### Record lookup mechanism

Like values, table records can also be inherited as a unit by multiple contexts, but they can also be partially redefined (*overwritten*), defined for a specific context (*root mode*), or be *occulted*.

Formally, a table record has one of four distinct definition modes:

|                                  |                                                                                                                                                                                                                                                                                                       |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b><i>root record</i></b>        | Locally defined in the table and has no parent. This means that no record with the same primary key exists in the parent table, or that this parent is an occulting record.                                                                                                                           |
| <b><i>overwriting record</i></b> | Locally defined in the table and has a parent record. This means that a record with the same primary key exists in the parent table, and that this parent is not an occulting record. The overwriting record inherits its values from its parent, except for the values that it explicitly redefines. |
| <b><i>inherited record</i></b>   | Not locally defined in the current table and has a parent record. All values are inherited.<br><br>Functions are always resolved in the current record context and are not inherited.                                                                                                                 |
| <b><i>occulting record</i></b>   | Specifies that, if a parent with the same primary key is defined, this parent will not be visible in table descendants.                                                                                                                                                                               |

See also [Dataset inheritance](#) [p 125]

### Defining inheritance behavior at the table level

It is also possible to specify management rules in the declaration of a table in the data model.

See also [Properties related to dataset inheritance](#) [p 499]

## 48.3 Inherited fields

The specific inheritance mechanism allows fetching a value of a field according to its relationship to other tables.

### Field inheritance declaration

Specific inheritance must be specified on terminal nodes in the underlying data model and is declared as follows:

```
<xs:element name="sampleInheritance" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:inheritance>
        <sourceRecord>
          /root/table1/fkTable2, /root/table2/fkTable3
        </sourceRecord>
      </osd:inheritance>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

```

    <sourceNode>color</sourceNode>
  </osd:inheritance>
</xs:appinfo>
</xs:annotation>
</xs:element>

```

The element `sourceRecord` is an expression that describes how to look up the record from which the value is inherited. It is a foreign key, or a sequence of foreign keys, from the current element to the source table.

If `sourceRecord` is not defined in the data model, the inherited fields are fetched from the current record.

The element `sourceNode` is the path of the node from which to inherit in the source record.

The following conditions must be satisfied for specific inheritance:

- The element `sourceNode` is mandatory.
- The expression for the path to the source record must be a consistent path of foreign keys, from the current element to the source record. This expression must involve only one-to-one and zero-to-one relationships.
- The `sourceRecord` cannot contain any aggregated list elements.
- Each element of the `sourceRecord` must be a foreign key.
- If the inherited field is also a foreign key, the `sourceRecord` cannot refer to itself to get the path to the source record of the inherited value.
- Every element of the `sourceRecord` must exist.
- The source node must belong to the table containing the source record.
- The source node must be terminal.
- The source node must be writeable.
- The source node type must be compatible with the current node type.
- The source node cardinalities must be compatible with those of the current node.
- The source node cannot be the same as the inherited field if the fields to inherit from are fetched into the same record.

### ***Value lookup mechanism***

The lookup mechanism for inherited fields values proceeds as follows:

1. If the value is locally defined, it is returned.  
It can be explicitly `null`
2. Otherwise, looks up the source record and value to inherit from, according to the properties that are defined in the data model.
3. The process is recursive; if the source node does not locally define a value, it is then looked up according to the inheritance behavior of the source node.

## **48.4 Optimize & Refactor service**

EBX® provides a built-in user service for optimizing the dataset inheritance in the hierarchy of datasets. This service performs the following functions:

- **Handles duplicated values:** Detects and removes all parameter values that are duplicates of the inherited value.
- **Mutualizes common values:** Detects and mutualizes the common values among the descendants of a common ancestor.

## Procedure details

Datasets are processed from the bottom up, which means that if the service is run on the dataset at level  $N$ , with  $N+1$  being the level of its children and  $N+2$  being the level of its children's children, the service will first process the datasets at level  $N+2$  to determine if they can be optimized with respect to the datasets at level  $N+1$ . Next, it would proceed with an optimization of level  $N+1$  against level  $N$ .

### Note

- These optimization and refactoring functions do not handle default values that are declared in the data model.
- The highest level considered during the optimization procedure is always the dataset on which the service is run. This means that optimization and refactoring are not performed between the target dataset and its own ancestors.
- Table optimization is performed on records with the same primary key.
- Inherited fields are not optimized.
- *The optimization and refactoring functions do not modify the resolved view of a dataset, if it is activated.*

## Service availability

The 'Optimize & Refactor' service is available on datasets that have child datasets and have the 'Activated' property set to 'No' in their dataset information.

The service is available to any profile with write access on current dataset values. It can be disabled by setting restrictive access rights on a profile.

### Note

For performance reasons, access rights are not verified on every node and table record.



## CHAPTER 49

# Permissions

Permissions dictate the access each user has to data and actions.

This chapter contains the following topics:

1. [Overview](#)
2. [Defining user-defined rules](#)
3. [Defining programmatic rules](#)
4. [Resolving permissions on data](#)
5. [Resolving permissions on services](#)
6. [Resolving permissions on actions](#)

## 49.1 Overview

Permissions are related to whether actions are authorized or not. They are also related to access rights, that is, whether an entity is hidden, read, or read-write. The main entities controlled by permissions are:

- Dataspace
- Dataset
- Table
- Group
- Field

### ***Users, roles and profiles***

The definition and resolution of permissions make extensive use of the notion of *profiles*, which is the generic term applied to users or roles.

Each user can participate in several roles, and a role can be shared by several users.

These relationships are defined in the user and roles directory. See [Users and roles directory](#) [p 401].

#### **Special definitions:**

- An *administrator* is a member of the built-in role 'ADMINISTRATOR'.
- An *owner of a dataset* is a member of the *owner* attribute specified in the information of a root dataset. In this case, the built-in role 'OWNER' is activated when permissions are resolved in the context of the dataset.

- An *owner of a dataspace* is a member of the *owner* attribute specified for a dataspace. In this case, the built-in role 'OWNER' is activated when permissions are resolved in the context of the dataspace.

## Permission rules

A permission rule defines the authorization granted to a profile for a particular entity.

User-defined permission rules are created through the user interface. See the section [Defining user-defined rules](#) [p 276].

Programmatic permission rules can be created by developers. See the section [Defining programmatic rules](#) [p 280].

## Resolution of permissions

Permissions are always resolved in the context of an authenticated user session, thus permissions are mainly based on the user profiles.

In general, resolution of permissions is performed restrictively between a given level and its parent level. Thus, at any given level, a user cannot have a higher permission than the one resolved at a parent level.

Programmatic permissions are always considered to be restrictive.

### Note

In the Java API, the class `SessionPermissionsAPI` provides access to the resolved permissions.

### See also

[Resolving permissions on data](#) [p 281]

[Resolving permissions on services](#) [p 285]

[Resolving permissions on actions](#) [p 287]

## Owner and administrator special permissions

### On a dataset

An administrator or owner of a dataset can perform the following actions:

- Manage its permissions
- Change its owner, if the dataset is a root dataset
- Change its general information (localized labels and descriptions)

### Attention

While the definition of permissions can restrict an administrator or dataset owner's right to view data or perform certain actions, it remains possible for them to modify their own access, as they will always have access to permissions management.

### On a dataspace

To be a *super owner* of a dataspace, a user must either:

- Own the dataspace and be allowed to manage its permissions, or
- Own a dataspace that is an ancestor of the current dataspace and be allowed to manage the permissions of that ancestor dataspace.

An administrator or super owner of a dataspace can perform the following actions:

- Manage its permissions of dataspace.
- Change its owner
- Lock it or unlock it
- Change its general information (localized labels and descriptions)

Furthermore, in a workflow, when using a "Create a dataspace" or "Create a snapshot" built-in script task, resolved permissions are computed using the owner defined in the script task's configuration, rather than the current session. This is because, in these cases, the current session is associated with a system user.

#### Attention

While the definition of permissions can restrict an administrator or dataspace owner's right to view data or perform certain actions, it remains possible for them to modify their own access, as they will always have access to permissions management.

### ***Impact of merge on permissions***

When a dataspace is merged, the permissions of the child dataset are merged with those of the parent dataspace if and only if the user specifies to do so during the merge process. The permissions of its parent dataspace are never impacted.

If some elements are hidden for the profile attempting to perform a merge, it will not be possible to proceed as the impacts of the merge on data will not be fully visible.

### ***Important considerations about permissions***

The following statements should be kept in mind while working with permissions:

- Whenever the hidden permission is returned for a session on a table child node, a user could "guess" sensitive information by filtering or sorting on this node in the following cases:
  - When using a Request<sup>API</sup> on this table with **permissions enabled** Request.setSession<sup>API</sup>. To avoid this, permissions on this node must be checked explicitly before applying the filter or the sort criteria.
  - When defining a custom view on this table in the UI. To avoid this, view definition permissions should be restricted for such users.
- Resolution of custom display labels for tables ('defaultLabel' property) and relationships ('display' property) ignores permission, and fields usually hidden due to access rights restrictions will be displayed in such labels. As a result, these labels should not contain any confidential field. Otherwise, a permission strategy should also be defined to restrict the display of the whole label.
- When a procedure disables all permission checks by using ProcedureContext.setAllPrivileges<sup>API</sup>, the client code must check that the current user session is allowed to run the procedure.

- When performing actions on a table (create, delete, overwrite or occult) in a procedure, the current user session access right on the table node is ignored during the permission resolution. Should this check be performed, the client code must explicitly call `SessionPermissions.getNodeAccessPermissionAPI` beforehand in the procedure.
- To optimize the resolution of permissions for both data and user services, a dedicated cache is implemented at the session level; it only takes user-defined permissions into account, not programmatic rules (which are not cached since they are contextual and dynamic). The session cache life cycle depends on the context, as described hereafter:
  - In the UI, the cache is cleared for every non-ajax event (i.e on page display, pop-up opening, etc.).
  - In programmatic procedures, the cache lasts until the end of the procedure, unless explicitly cleared (see below).

**Attention**

When modifying permissions in a procedure context (by importing an EBX® archive or merging a dataspace programmatically), the session cache **must** be cleared via a call to `Session.clearCacheAPI`. Otherwise, these modifications will not be reflected until the end of the procedure.

## 49.2 Defining user-defined rules

Each level has a similar schema, which allows defining permission rules for profiles.

## Defining dataspace user-defined rules

For a given dataspace, the allowable permissions for each profile are as follows:

| Dataspace access | Authorization                                                                                                                                                                                                                                                                                                                        |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Write            | <ul style="list-style-type: none"> <li>• Can view the dataspace.</li> <li>• Can access datasets according to dataset permissions.</li> </ul>                                                                                                                                                                                         |
| Read-only        | <ul style="list-style-type: none"> <li>• Can view the dataspace and its snapshots.</li> <li>• Can view child dataspace, if allowed by permissions.</li> <li>• Can view contents of the dataspace, though cannot modify them.</li> </ul>                                                                                              |
| Hidden           | <ul style="list-style-type: none"> <li>• Can neither see the dataspace nor its snapshots.</li> <li>• If allowed to view child dataspace, can see the current dataspace but cannot select it.</li> <li>• Cannot access the dataspace contents, including datasets.</li> <li>• Cannot perform any actions on the dataspace.</li> </ul> |

|                                 |                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Restriction policy</b>       | Indicates whether this dataspace profile-permission association should have priority over other permissions rules.         |
| <b>Create a child dataspace</b> | Indicates whether the profile can create child dataspace from the current dataspace.                                       |
| <b>Create a child snapshot</b>  | Indicates whether the profile can create snapshots of the current dataspace.                                               |
| <b>Initiate merge</b>           | Indicates whether the profile can merge the current dataspace with its parent dataspace.                                   |
| <b>Export archive</b>           | Indicates whether the profile can export the current dataspace as an archive.                                              |
| <b>Import archive</b>           | Indicates whether the profile can import an archive into the current dataspace.                                            |
| <b>Close a dataspace</b>        | Indicates whether the profile can close the current dataspace.                                                             |
| <b>Close a snapshot</b>         | Indicates whether the profile can close a snapshot of the current dataspace.                                               |
| <b>Rights on services</b>       | Indicates if a profile has the right to execute services on the dataspace. By default, all dataspace services are allowed. |

An administrator or super owner of the current dataspace or a given user who is allowed to modify permissions on the current dataspace can modify these permissions to restrict dataspace services for certain profiles.

---

**Permissions of child dataspace when created**

When a user creates a child dataspace, the permissions of this new dataspace are automatically assigned to the profile's owner, based on the permissions defined under 'Permissions of child dataspace when created' in the parent dataspace. If multiple permissions are defined for the owner through different roles, the owner's profile behaves like any other profile and [permissions are resolved](#) [p 274] as usual.

---

## ***Defining dataset user-defined rules***

For a given dataset, the allowable permissions for each profile are as follows:

### **Actions on datasets**

---

|                                  |                                                                                                                  |
|----------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Restriction policy</b>        | Indicates whether this dataset profile-permission association should have priority over other permissions rules. |
| <b>Create a child dataset</b>    | Indicates whether the profile has the right to create a child dataset of the current dataset.                    |
| <b>Duplicate dataset</b>         | Indicates whether the profile has the right to duplicate the current dataset.                                    |
| <b>Change the dataset parent</b> | Indicates whether the profile has the right to change the parent dataset of a given child dataset.               |

---

## Actions on tables

The action rights on default tables are defined at the dataset level. It is then possible to override these default rights for one or more tables. The allowable permissions for each profile are as follows:

|                                   |                                                                                          |
|-----------------------------------|------------------------------------------------------------------------------------------|
| <b>Create a new record</b>        | Indicates whether the profile has the right to create records in the table.              |
| <b>Overwrite inherited record</b> | Indicates whether the profile has the right to overwrite inherited records in the table. |
| <b>Occult inherited record</b>    | Indicates whether the profile has the right to occult inherited records in the table.    |
| <b>Delete a record</b>            | Indicates whether the profile has the right to delete records in the table.              |

## Access rights on node values

Permissions defined on specific terminal nodes override their default access rights.

|                   |                                                 |
|-------------------|-------------------------------------------------|
| <b>Read-write</b> | Can view and modify node values.                |
| <b>Read</b>       | Can view nodes, but cannot modify their values. |
| <b>Hidden</b>     | Cannot view nodes.                              |

## Permissions on services

An administrator or an owner of the current dataspace can modify the service default permission to either restrict or grant access to certain profiles.

|                 |                                                                                                                                                                                                                                     |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enabled</b>  | Grants service access to the current profile.                                                                                                                                                                                       |
| <b>Disabled</b> | Forbids service access to the current profile. It will not be displayed in menus, nor will it be launchable via web components.                                                                                                     |
| <b>Default</b>  | <p>Sets the service permission to enabled or disabled, according to the default permission defined upon service declaration.</p> <p>See <code>ActivationContext.setDefaultPermission<sup>API</sup></code> for more information.</p> |

## 49.3 Defining programmatic rules

Programmatic rules give the possibility to define more precisely the conditions for accessing data or user services depending on the context.

There are different types of programmatic rules:

- the `AccessRuleAPI`, described in the section below [Defining access rules on data](#) [p 280].
- the [ServiceActivationRule](#) [p 280], described in the section below [Defining activation rules on service](#) [p 280].
- the `ServicePermissionRuleAPI`, described in the section below [Defining permission rules on service](#) [p 281].

### **Defining access rules on data**

`AccessRules` are rules that programmatically define, depending on the context, the read/write rights on a data model node or on the records of a table.

The definition of an `AccessRule` is performed as follows:

1. Creation of a rule in the form of a Java class implementing the `AccessRuleAPI` or `AccessRuleForCreateAPI` interface.
2. Assignment of this rule to concerned nodes in the schema extension: `SchemaExtensionsAPI`.

According to the rule target (model node(s) or records) and type (`AccessRule` or `AccessRuleForCreate`), several methods such as `SchemaExtensionsContext.setAccessRuleForCreateOnNodeAPI` or `SchemaExtensionsContext.setAccessRuleOnOccurrenceAPI` can be used.

The rule thus assigned is said to be "local" and is only executed when the target entity is requested. See [Resolving permissions on data](#) [p 281] for more information.

#### **Attention**

Only one `AccessRule` can be defined for each node, dataspace or record. Only one `AccessRuleForCreate` can be defined for each table child node. The definition of a new programmatic rule of one type will lead to the replacement of the existing one.

### **Defining activation rules on service**

The `ServiceActivationRules` allow to specify if a service is activated or not for a given dataspace or dataset. A service that has been deactivated through this rule is never available in the entity for which it is deactivated, regardless of the current profile, for execution or display, even in permission screens.

The definition of a `ServiceActivationRule` is carried out as follows:

1. Creation of a rule in the form of a Java class implementing the `ServiceActivationRuleForDataspaceAPI` interface or `ServiceActivationRuleForDatasetAPI`, depending on the service type.
2. Assignment of this rule to the impacted services at their declaration level, depending on the service type, via the `ActivationContextOnDataspace.setActivationRuleAPI` or `ActivationContextWithDatasetSet.setActivationRuleAPI` methods.



The resulting assigned rule will be evaluated during the service activation evaluation. See [Resolving permissions on services](#) [p 285] for more information.

## Defining permission rules on service

The `ServicePermissionRules` are advanced rules allowing to dynamically define the display and execution conditions of a service depending on the context (current session, selected entity, etc.). The service should be activated for the current context beforehand for this type of rule to be triggered.

The definition of a `ServicePermissionRule` is carried out as follows:

1. Creation of a rule in the form of a Java class implementing the `ServicePermissionRule`<sup>API</sup> interface.
2. Assignment of this rule to the impacted services:
  - Either, for new services, at their declaration level via the `ActivationContext.setPermissionRule`<sup>API</sup> method.

The rule thus assigned is said to be "global" and is only executed when the service is activated for the current context. See [Resolving permissions on services](#) [p 285] for more information.

- Or, for existing services, in the **schema extension** `SchemaExtensions`<sup>API</sup> via the `SchemaExtensionsContext.setServicePermissionRuleOnNode`<sup>API</sup> and `SchemaExtensionsContext.setServicePermissionRuleOnNodeAndAllDescendants`<sup>API</sup> methods. It is thus possible to assign a rule to any service, including standard services provided by EBX®, on one or more data model nodes: a table node, an association node, etc. The rule thus assigned is said to be "local" and is only executed in the extended schema context and when the node corresponds to the one specified. See [Resolving permissions on services](#) [p 285] for more information.

### Attention

Only one `ServicePermissionRule` can be defined for each model node. Thus, the definition of a new programmatic rule will replace the existing one.

## 49.4 Resolving permissions on data

### Resolving user-defined rules

Access rights defined using the user interface are resolved on four levels: dataspace, dataset, record (if applicable) and node.

If a profile is associated with restrictive access rights at a given level, the minimum of all restrictive rights defined at that level is resolved. If no restrictions are defined at that level, the maximum of all access rights defined at that level is resolved.

When a restrictive permission is defined for a profile, it takes precedence over the other permissions potentially granted by the user's other roles. Generally, for all user-defined permission rules that match the current user session:

- If some rules with restrictions are defined, the minimum permissions of these restricted rules are applied.

- If no rules having restrictions are defined, the maximum permissions of all matching rules are applied.

**Examples:**

Given two profiles *P1* and *P2* concerning the same user, the following table lists the possibilities when resolving that user's permission to a service.

| P1 authorization | P2 authorization | Permission resolution                                |
|------------------|------------------|------------------------------------------------------|
| Enabled          | Enabled          | Enabled. Restrictions do not make any difference.    |
| Disabled         | Disabled         | Disabled. Restrictions do not make any difference.   |
| Enabled          | Disabled         | Enabled, unless P2's authorization is a restriction. |
| Disabled         | Enabled          | Enabled, unless P1's authorization is a restriction. |

The same restriction policy is applied for data access rights resolution.

In another example, a dataspace can be hidden from all users by defining a restrictive association between the built-in profile "Profile.EVERYONE" and the access right "hidden".

At any given level, the most restrictive access rights between those resolved at this level and higher levels are applied. For instance, if a user's dataset access permissions resolve to read-write access, but the container dataspace only allows read access, the user will only have read-only access to this dataset.

**Note**

The dataset inheritance mechanism applies to both values and access rights. That is, access rights defined on a dataset will be applied to its child datasets. It is possible to override these rights in the child dataset.

## Access rights resolution example

In this example, there are three users who belong to the following defined roles and profiles:

| User          | Profile                                                                                                         |
|---------------|-----------------------------------------------------------------------------------------------------------------|
| <b>User 1</b> | <ul style="list-style-type: none"> <li>• user1</li> <li>• role A</li> <li>• role B</li> </ul>                   |
| <b>User 2</b> | <ul style="list-style-type: none"> <li>• user2</li> <li>• role A</li> <li>• role B</li> <li>• role C</li> </ul> |
| <b>User 3</b> | <ul style="list-style-type: none"> <li>• user3</li> <li>• role A</li> <li>• role C</li> </ul>                   |

The access rights of the profiles on a given element are as follows:

| Profile | Access rights | Restriction policy |
|---------|---------------|--------------------|
| user1   | Hidden        | Yes                |
| user3   | Read          | No                 |
| Role A  | Read/Write    | No                 |
| Role B  | Read          | Yes                |
| Role C  | Hidden        | No                 |

After resolution based on the role and profile access rights above, the rights that are applied to each user are as follows:

| User          | Resolved access rights |
|---------------|------------------------|
| <b>User 1</b> | Hidden                 |
| <b>User 2</b> | Read                   |
| <b>User 3</b> | Read/Write             |

## Resolving dataspace and snapshot access rights

At dataspace level, access rights are resolved as follows:

- If a user has several rights defined through multiple profiles:
  - If the rights include restrictions, the minimum of the restrictive profile-rights associations is applied.
  - Otherwise, the maximum of the profile-rights associations is applied.
- If the user has no rights defined:
  - If the user is an administrator or owner of the dataspace, read-write access is given for this dataspace.
  - Otherwise, the dataspace will be hidden.

## Resolving dataset access rights

At the dataset level, the same principle applies as at the dataspace level. After resolving the access rights at the dataset level alone, the final access rights are determined by taking the minimum rights between the resolved dataspace rights and the resolved dataset rights. For example, if a dataspace is resolved to be read-only for a user and one of its datasets is resolved to be read-write, the user will only have read-only access to that dataset.

## Resolving node access rights

At the node level, the same principle applies as at the dataspace and dataset levels. After resolving the access rights at the node level alone, the final access rights are determined by taking the minimum rights between the resolved dataset rights and the resolved node rights.

Specific access rights can be defined at the node level. If no specific access right is defined, the default access right is used for the resolution process.

### Note

The resolution procedure is slightly different for table and table child nodes.

## Special case for table and table child nodes

This describes the resolution process used for a given table node or table record  $N$ .

For each user-defined permission rule that matches one of the user's profiles, the access rights for  $N$  are either:

1. The locally defined access rights for  $N$ ;
2. Inherited from the access rights defined on the table node;
3. Inherited from the default access rights for dataset values.

All matching user-defined permission rules are used to resolve the access rights for  $N$ . Resolution is done according to the [restriction policy](#) [p 281].

The final resolved access rights will be the minimum between the dataspace, dataset and the resolved access right for  $N$ .

## Resolving programmatic rules

There are three levels of resolution for programmatic access right rules: dataset, record and node. Since only one programmatic access rule can be set for a given level, the last rule set is the one used by the resolution procedure.

### Rule resolution on dataset

For a dataset, the last rule set is considered as the resolved rule

### Rule resolution on record

For a record, the resolved rule is the minimum between the resolved rule set on the dataset and the rule set on this record.

See `SchemaExtensionsContext.setAccessRuleOnOccurrenceAPI` for more details.

### Rule resolution on node

For a node that is a child node of a record, the resolved rule is the minimum between the resolved rule on the record and the rule set on this node.

For a child node of a dataset, the resolved rule is the minimum between the resolved rule set on the dataset and the rule set on this node.

See `SchemaExtensionsContext.setAccessRuleOnNodeAPI` for more details.

### Display policy for foreign key drop-down menus

If a record is hidden due to access rules, it will not appear in foreign key drop-down menus.

#### Attention

The resolved access rights on a dataset or dataset node is the minimum between the resolved access rights defined in the user interface and the resolved programmatic rules, if any.

## 49.5 Resolving permissions on services

User services give the possibility to execute specific and advanced features from the user interface. Depending on their definition, these services can be called from a menu, as an action in a workflow, as a perspective item, or can be executed directly from a URL as a [Web component](#) [p 194].

See also [Overview](#) [p 565]

The permissions of a service are resolved as the service is called from the user interface, namely:

- During the execution, just before the service is displayed.  
If the permission resolved in the user context is not enabled, a restriction message is displayed in place of the service.
- During the display of menus if the service is defined as displayable in menus.  
If the permission resolved in the context for the user is not enabled, the service will not be displayed in the menu.

Thus, upon every request the resolution of permissions for a service is carried out as follows, in the following order and as long as conditions are respected:

1. The service activation has to correspond to the current context. This activation considers:
  - the selected entity type (dataset, table, record, etc.);
  - static activation rules defined within the `UserServiceDeclaration.defineActivationAPI` method;
  - the potential dynamic activation rule ([ServiceActivationRule](#) [p 280]) also defined within the `UserServiceDeclaration.defineActivationAPI` method.
2. When the service is activated for the current context, permissions for the user session will be evaluated:
  - If permissions have been defined via the user interface for the current user (or for their roles), their resolution must return enabled.

For more information, please refer to the [Resolving user-defined rules](#) [p 286] section.

  - If a [global permission rule](#) [p 281] is defined for the service, it must return enabled for the context provided (see `ServicePermissionRuleContextAPI`).
  - If a [local permission rule](#) [p 281] is defined for the selected node, it must return enabled for the context provided (see `ServicePermissionRuleContextAPI`).

## Resolving user-defined rules

### Example

In this example, there are two users belonging to different roles and profiles:

| User   | Profiles                                                                                      |
|--------|-----------------------------------------------------------------------------------------------|
| User 1 | <ul style="list-style-type: none"> <li>• user1</li> <li>• role A</li> <li>• role B</li> </ul> |
| User 2 | <ul style="list-style-type: none"> <li>• role C</li> <li>• role D</li> </ul>                  |

The permissions associated with the roles and profiles defined on the dataset level are as follows:

| Profile | Built-in service create (@creation) | Built-in service duplicate (@duplicate) | Built-in service compare (@compare) | Custom service 1 (custom1) | Custom service 2 (custom2) | Restriction policy |
|---------|-------------------------------------|-----------------------------------------|-------------------------------------|----------------------------|----------------------------|--------------------|
| user1   | Enabled                             | Disabled                                | Enabled                             | Disabled                   | Enabled                    | No                 |
| Role A  | Enabled                             | Enabled                                 | Disabled                            | Enabled                    | Disabled                   | Yes                |
| Role B  | Enabled                             | Disabled                                | Enabled                             | Enabled                    | Disabled                   | Yes                |
| Role C  | Enabled                             | Enabled                                 | Disabled                            | Disabled                   | Disabled                   | No                 |
| Role D  | Enabled                             | Disabled                                | Disabled                            | Enabled                    | Disabled                   | No                 |

The services available to each user after permission resolution are as follows:

| Users  | Available services                      |
|--------|-----------------------------------------|
| User 1 | Built-in service create (@creation)     |
|        | Custom service 1 (custom1)              |
| User 2 | Built-in service create (@creation)     |
|        | Built-in service duplicate (@duplicate) |
|        | Custom service 1 (custom1)              |

See also [Resolving user-defined rules](#) [p 281]

## 49.6 Resolving permissions on actions

Actions are low-level operations for EBX® object manipulation on which it is possible to define execution rights for a profile. Unlike permissions on user services, which only impact the user interface, these rights are also applicable when an operation is carried out programmatically (i.e. via a Procedure<sup>API</sup>) or indirectly (for example during data import, actions on the table (create, override, occult and delete) are evaluated).

Here is the list of actions on which rights can be defined:

| Action object    | Available actions               |
|------------------|---------------------------------|
| <b>Dataspace</b> | Create a child dataspace        |
|                  | Create a snapshot               |
|                  | Launch a merge                  |
|                  | Export an archive               |
|                  | Import an archive               |
|                  | Close the dataspace             |
|                  | Close the snapshot              |
|                  | Create a dataset                |
| <b>Dataset</b>   | Duplicate the dataset           |
|                  | Delete the dataset              |
|                  | Activate/deactivate the dataset |
|                  | Create a view                   |
| <b>Table</b>     | Create a new record             |
|                  | Override records                |
|                  | Occult records                  |
|                  | Delete records                  |

For the resolution of permissions on actions, only the permissions defined via the user interface for the current user (or their roles) will be taken into account, the restriction policy being applied as for any other permission defined via the user interface.

For more information, please refer to the [Resolving user-defined rules](#) [p 289] section.



## Resolving user-defined rules

### Example

In this example, we have two users belonging to different roles and profiles:

| User   | Profiles                                                                                      |
|--------|-----------------------------------------------------------------------------------------------|
| User 1 | <ul style="list-style-type: none"> <li>• user1</li> <li>• role A</li> <li>• role B</li> </ul> |
| User 2 | <ul style="list-style-type: none"> <li>• role C</li> <li>• role D</li> </ul>                  |

Rights associated with roles and profiles on the actions of a given table are as follows:

| Profile | Create a record | Override a record | Occult a record | Delete a record | Restriction policy |
|---------|-----------------|-------------------|-----------------|-----------------|--------------------|
| user1   | No              | Yes               | No              | Yes             | No                 |
| Role A  | Yes             | No                | Yes             | No              | Yes                |
| Role B  | No              | Yes               | Yes             | No              | Yes                |
| Role C  | Yes             | No                | No              | No              | No                 |
| Role D  | No              | No                | Yes             | No              | No                 |

The actions available to each user after resolving the rights are as follows:

| Users  | Available actions |
|--------|-------------------|
| User 1 | Occult a record   |
| User 2 | Create a record   |
|        | Occult a record   |

See also [Resolving user-defined rules](#) [p 281]



## CHAPTER 50

---

# Criteria editor

This chapter contains the following topics:

1. [Overview](#)
2. [Conditional blocks](#)
3. [Atomic criteria](#)

## 50.1 Overview

The criteria editor is included in several different areas of the user interface. It allows defining table filters, as well as validation and computation rules on data. This editor is based on the XPath 1.0 W3C Recommendation.

Two types of criteria exist: atomic criteria and conditional blocks.

See also [Supported XPath syntax](#) [p 225]

## 50.2 Conditional blocks

Conditional blocks are made up of atomic criteria and other conditional blocks. They express a condition based on the criteria. The following types of blocks exist:

- **No criteria match:** None of the criteria in the block match.
- **Not all criteria match:** At least one criterion in the block does not match.
- **All criteria match:** All criteria in the block match.
- **At least one criterion matches:** One or more of the criteria match.

## 50.3 Atomic criteria

An atomic predicate is defined by a field, an operator, and an expression (either a value or an XPath formula).

|                  |                                                                                                                       |
|------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>Field</b>     | Specifies the field of the table to which the criterion applies.                                                      |
| <b>Operator</b>  | Specifies the operator used. Available operators depend on the data type of the field.                                |
| <b>Value</b>     | Specifies the value or expression. See <a href="#">Expression</a> [p 292] below.                                      |
| <b>Code only</b> | If checked, specifies searching the underlying values for the field instead of labels, which are searched by default. |

### ***Expression***

The expression can either be a fixed value or a formula. When creating a filter, only fixed values are authorized. During creation of a validation or computation rule, a formula can be created using the wizard.

**Known limitation:** The formula field does not validate input values, only the syntax and path are checked.

## CHAPTER 51

# Performance guidelines

This chapter contains the following topics:

1. [Basic performance checklist](#)
2. [Checklist for dataspace usage](#)
3. [Memory management](#)
4. [Validation](#)
5. [Mass updates](#)
6. [Accessing tables](#)

## 51.1 Basic performance checklist

While TIBCO EBX® is designed to support large volumes of data, several common factors can lead to poor performance. Addressing the key points discussed in this section will solve the usual performance bottlenecks.

### *Expensive programmatic extensions*

For reference, the table below details the programmatic extensions that can be implemented.

| Use case             | Programmatic extensions that can be involved                                                                                                                                                                                                                                                                     |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Validation           | <ul style="list-style-type: none"> <li>• <b>programmatic constraints</b> <code>Constraint<sup>API</sup></code></li> <li>• <b>computed values</b> <code>ValueFunction<sup>API</sup></code></li> </ul>                                                                                                             |
| Table access         | <ul style="list-style-type: none"> <li>• <b>record-level permission rules</b> <code>SchemaExtensionsContext.setAccessRuleOnOccurrence<sup>API</sup></code></li> <li>• <b>programmatic filters</b> <code>AdaptationFilter<sup>API</sup></code></li> </ul>                                                         |
| EBX® content display | <ul style="list-style-type: none"> <li>• <b>computed values</b> <code>ValueFunction<sup>API</sup></code></li> <li>• <b>UI Components</b> <code>UIBeanEditor<sup>API</sup></code></li> <li>• <b>node-level permission rules</b> <code>SchemaExtensionsContext.setAccessRuleOnNode<sup>API</sup></code></li> </ul> |
| Data update          | <ul style="list-style-type: none"> <li>• <b>triggers</b> <code>Package com.orchestranetworks.schema.trigger<sup>API</sup></code></li> </ul>                                                                                                                                                                      |

For large volumes of data, cumbersome algorithms have a serious impact on performance. For example, a constraint algorithm's complexity is  $O(n^2)$ . If the data size is 100, the resulting cost is

proportional to 10 000 (this generally produces an immediate result). However, if the data size is 10 000, the resulting cost will be proportional to 10 000 000.

Another reason for slow performance is calling external resources. Local caching usually solves this type of problem.

If one of the use cases above displays poor performance, it is recommended to track the problem either through code analysis or using a Java profiling tool.

## **Directory integration**

Authentication and permissions management involve the [user and roles directory](#) [p 401].

If a specific directory implementation is deployed and accesses an external directory, it can be useful to ensure that local caching is performed. In particular, one of the most frequently called methods is `Directory.isUserInRoleAPI`.

## **Aggregated lists**

In a data model, when an element's cardinality constraint `maxOccurs` is greater than 1 and no `osd:table` is declared on this element, it is implemented as a Java `List`. This type of element is called an [aggregated list](#) [p 492], as opposed to a table.

It is important to consider that there is no specific optimization when accessing aggregated lists in terms of iterations, user interface display, etc. Besides performance concerns, aggregated lists are limited with regard to many functionalities that are supported by tables. See [tables introduction](#) [p 495] for a list of these features.

### **Attention**

For the reasons stated above, aggregated lists should be used only for small volumes of simple data (one or two dozen records), with no advanced requirements for their identification, lookups, permissions, etc. For larger volumes of data (or more advanced functionalities), it is recommended to use `osd:table` declarations.

## **51.2 Checklist for dataspace usage**

[Dataspaces](#) [p 88] available in semantic mode, are an invaluable tool for managing complex data life cycles. While this feature brings great flexibility, it also implies a certain overhead cost, which should be taken into consideration for optimizing usage patterns.

This section reviews the most common performance issues that can appear in case of an intensive use of many dataspace containing large tables, and how to avoid them.

#### Note

Sometimes, the use of dataspace is not strictly needed. As an extreme example, consider the case where every transaction triggers the following actions:

1. A dataspace is created.
2. The transaction modifies some data.
3. The dataspace is merged, closed, then deleted.

In this case, no future references to the dataspace are needed, so using it to make isolated data modifications is unnecessary. Thus, using `ProcedureAPI` already provides sufficient isolation to avoid conflicts from concurrent operations. It would then be more efficient to directly do the modifications in the target dataspace, and get rid of the steps which concern branching and merging.

For a developer-friendly analogy, referring to a source-code management tool (CVS, SVN, etc.): when you need to perform a simple modification impacting only a few files, it is probably sufficient to do so directly on the main branch. In fact, it would be neither practical nor sustainable, with regard to file tagging/copying, if every file modification involved branching the whole project, modifying the files, then merging the dedicated branch.

## Insufficient memory

When a table is in semantic mode (default), the EBX® Java memory cache is used. It ensures a much more efficient access to data when this data is already loaded in the cache. However, if there is not enough space for working data, swaps between the Java heap space and the underlying database can heavily degrade overall performance.

This memory swap overhead can only occur for tables in a dataspace with an [on-demand loading strategy](#) [p 297].

Such an issue can be detected by looking at the [monitoring log file](#) [p 297]. If it occurs, various actions can be considered:

- reducing the number of child dataspace that contain large tables;
- reducing the number of indexes specifically defined for large tables;
- using relational mode instead of semantic mode;
- or (obviously) allocating more memory, or optimizing the memory used by applications for non-EBX® objects.

#### See also

[Memory management](#) [p 296]

[Relational mode](#) [p 243]

## Transaction cancels

In semantic mode, when a transaction has performed some updates in the current dataspace and then aborts, loaded indexes of the modified tables are reset. If updates on a large table are often cancelled and, at the same time, this table is intensively accessed, then the work related to index rebuild will

slow down the access to the table; moreover, the induced memory allocation and garbage collection can reduce the overall performance.

**See also**

**Functional guard and exceptions** `TableTrigger.guardAndException`<sup>API</sup>

*Procedure*<sup>API</sup>

## **Reorganization of database tables**

As with any database, inserting and deleting large volumes of data may lead to fragmented data, which can deteriorate performance over time. To resolve the issue, reorganizing the impacted database tables is necessary. See [Monitoring and cleanup of the relational database](#) [p 379].

A specificity of EBX® is that creating dataspace and snapshots adds new entries to tables HTA and ATB. When poor performance is experienced, it may be necessary to schedule a reorganization of these tables, for large repositories in which many dataspace are created and deleted.

**See also** [Monitoring and cleanup of the relational database](#) [p 379]

## **51.3 Memory management**

### **Loading strategy**

The administrator can specify the loading strategy of a dataspace or snapshot in its information. The default strategy is to load and unload the resources on demand. For resources that are heavily used, a *forced load* strategy is usually recommended.



The following table details the loading modes which are available in semantic mode. Note that the application server must be restarted so as to take into account any loading strategy change.

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>On-demand loading and unloading</b>  | <p>In this default mode, each resource in a dataspace is loaded or built only when it is needed. The resources of the dataspace are "soft"-referenced using the standard Java <code>SoftReference</code> class. This implies that each resource can be unloaded "at the discretion of the garbage collector in response to memory demand".</p> <p>The main advantage of this mode is the ability to free memory when needed. As a counterpart, this implies a load/build cost when an accessed resource has not yet been loaded since the server started up, or if it has been unloaded since.</p> |
| <b>Forced loading</b>                   | <p>If the forced loading strategy is enabled for a dataspace or snapshot, its resources are loaded asynchronously at server startup. Each resource of the dataspace is maintained in memory until the server is shut down or the dataspace is closed.</p> <p>This mode is particularly recommended for long-living dataspaces and/or those that are used heavily, namely any dataspace that serves as a reference.</p>                                                                                                                                                                             |
| <b>Forced loading and prevalidation</b> | <p>This strategy is similar to the forced loading strategy, except that the content of the loaded dataspace or snapshot will also be validated upon server startup.</p>                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Monitoring

Indications of EBX® load activity are provided by monitoring the underlying database, and also by the ['monitoring' logging category](#) [p 354].

If the numbers for *cleared* and *built* objects remain high for a long time, this is an indication that EBX® is swapping.

## Tuning memory

The maximum size of the memory allocation pool is usually specified using the Java command-line option `-xmx`. As is the case for any intensive process, it is important that the size specified by this option does not exceed the available physical RAM, so that the Java process does not swap to disk at the operating-system level.

Tuning the garbage collector can also benefit overall performance. This tuning should be adapted to the use case and specific Java Runtime Environment used.

## 51.4 Validation

The internal incremental validation framework will optimize the work required when updates occur. The incremental validation process behaves as follows:

- The first call to a dataset validation report performs a full validation of the dataset. The [loading strategy](#) [p 296] can also specify a dataspace to be prevalidated at server startup.
- Data updates will transparently and asynchronously maintain the validation report, insofar as the updated nodes specify explicit dependencies. For example, standard and static facets, foreign key constraints, dynamics facets, selection nodes specify explicit dependencies.
- If a mass update is executed or if there are too many validation messages, the incremental validation process is stopped. The next call to the validation report will then trigger a full validation.
- If a transaction is cancelled, the validation state of the updated dataset is reset. The next call to the validation report will trigger a full validation as well.

Certain nodes are systematically revalidated, however, even if no updates have occurred since the last validation. These are the nodes with *unknown dependencies*. A node has unknown dependencies if:

- It specifies a **programmatic constraint** `ConstraintAPI` in the default *unknown dependencies* mode,
- It declares a **computed value** `ValueFunctionAPI`, or it declares a dynamic facet that depends on a node that is itself a **computed value** `ValueFunctionAPI`.
- It is an [Inherited fields](#) [p 270] or it declares a dynamic facet that depends on a node that is itself an [Inherited fields](#) [p 270].

Consequently, on large tables (beyond the order of  $10^5$ ), it is recommended to avoid nodes with unknown dependencies (or at least to minimize the number of such nodes). For programmatic constraints, the developer is able to specify two alternative modes that drastically reduce incremental validation cost: *local dependency* mode and *explicit dependencies*. For more information, see **Dependencies and validation** `DependenciesDefinitionContext.dependenciesAPI`.

### Note

It is possible for an administrator user to manually reset the validation report of a dataset. This option is available from the validation report section in EBX®.

## 51.5 Mass updates

Mass updates can involve several hundred thousands of insertions, modifications and deletions. These updates are usually infrequent (usually initial data imports), or are performed non-interactively (nightly batches). Thus, performance for these updates is less critical than for frequent or interactive operations. However, similar to classic batch processing, it has certain specific issues.

### Batch mode

For relational tables, the implementation of insertions, updates and deletions relies on the JDBC batch feature. On large procedures, this can dramatically improve performance by reducing the number of round-trips between the application server and the database engine.

In order to fully exploit this feature, the batch mode can be activated on large procedures. See `ProcedureContext.setBatchAPI`. This disables the explicit check for existence before record insertions,

thus reducing the number of queries to the database, and making the batch processing even more efficient.

## Transaction boundaries

It is generally not advised to use a single transaction when the number of atomic updates in the transaction is beyond the order of  $10^4$ . Large transactions require a lot of resources, in particular, memory, from EBX® and from the underlying database.

To reduce transaction size, it is possible to:

- Specify the property [ebx.manager.import.commit.threshold](#) [p 363]. However, this property is only used for interactive archive imports performed from the EBX® user interface.
- Explicitly specify a **commit threshold** `ProcedureContext.setCommitThresholdAPI` inside the batch procedure.
- Structurally limit the transaction scope by implementing `ProcedureAPI` for a part of the task and executing it as many times as necessary.

On the other hand, specifying a very small transaction size can also hinder performance, due to the persistent tasks that need to be done for each commit.

### Note

If intermediate commits are a problem because transactional atomicity is no longer guaranteed, it is recommended to execute the mass update inside a dedicated dataspace. This dataspace will be created just before the mass update. If the update does not complete successfully, the dataspace must be closed, and the update reattempted after correcting the reason for the initial failure. If it succeeds, the dataspace can be safely merged into the original dataspace.

## Triggers

If required, triggers can be deactivated using the method `ProcedureContext.setTriggerActivationAPI`.

# 51.6 Accessing tables

## Functionalities

Tables are commonly accessed through EBX® and also through the `RequestAPI` API and data services. This access involves a unique set of functions, including a *dynamic resolution* process. This process behaves as follows:

- **Inheritance:** Inheritance in the dataset tree takes into account records and values that are defined in the parent dataset, using a recursive process. Also, in a root dataset, a record can inherit some of its values from the data model default values, defined by the `xs:default` attribute.
- **Value computation:** A node declared as an `osd:function` is always computed on the fly when the value is accessed. See `ValueFunction.getValueAPI`.
- **Filtering:** An [XPath predicate](#) [p 225], a **programmatic filter** `AdaptationFilterAPI`, or a record-level **permission rule** `SchemaExtensionsContext.setAccessRuleOnOccurrenceAPI` requires a selection of records.
- **Sort:** A sort of the resulting records can be performed.

## Accessing tables in semantic mode

### Architecture and design

In order to improve the speed of operations on tables, indexes are managed by the EBX® engine.

EBX® advanced features, such as advanced life-cycle (snapshots and dataspaces), dataset inheritance, and flexible XML Schema modeling, have led to a specialized design for indexing mechanisms. This design can be summarized as follows:

- *Indexes* maintain an in-memory data structure on a whole table.
- An index is not persisted, and building it requires loading all table blocks from the database.

#### Attention

Faster access to tables is ensured if indexes are ready and maintained in memory cache. As mentioned above, it is important for the Java Virtual Machine to have enough space allocated, so that it does not release indexes too quickly.

### Performance considerations

The request optimizer favors the use of indexes when computing a request result.

#### Attention

- Only XPath filters are taken into account for index optimization.
- Non-primary-key indexes are not taken into account for child datasets.

Assuming the indexes are already built, the impacts on performance are as follows:

1. If the request does not involve filtering, programmatic rules, or sorting, accessing its first few rows (these fetched by a paged view) is almost instantaneous.
2. If the request can be resolved without an extra sort step (this is the case if it has no sort criteria, or if its sort criteria relate to those of the index used for computing the request), accessing the first few rows of a table should be fast. More precisely, it depends on the cost of the specific filtering algorithm that is executed when fetching at least 2000 records.
3. Both cases above guarantee an access time that is independent of the size of the table, and provide a view sorted by the index used. If an extra sort is required, the time taken by the first access depends on the table size according to an  $N\log(N)$  function, where  $N$  is the number of records in the resolved view.

#### Note

The paginated requests automatically add the primary key to the end of the specified criterion, in order to ensure consistent ordering. Thus, the primary key fields should also be added to the end of any index intended to improve the performance of paginated requests. These include tabular and hierarchical views, and drop-down menus for table references.

If indexes are not yet built, or have been unloaded, additional time is required. The build time is  $O(N\log(N))$ .

Accessing the table data blocks is required when the request cannot be computed against a single index (whether for resolving a rule, filter or sort), as well as for building the index. If the table blocks are not present in memory, additional time is needed to fetch them from the database.

It is possible to get information through the [monitoring](#) [p 297] and request logging categories.

## Other operations on tables

The new records creations or record insertions depend on the primary key index. Thus, a creation becomes almost immediate if this index is already loaded.

## REST access to history table

The merge information in history table (the `merge_info` field) has a potentially high access cost. To improve performance and if the client code does not need this field, the [includeMergeInfo](#) [p 675] parameter must be set to `false`.

See [History](#) [p 249] for more information.

## Accessing tables in relational mode

When computing a request result, the EBX® engine delegates the following to the RDBMS:

- Handling of all request sort criteria, by translating them to an `ORDER BY` clause.
- Whenever possible, handling of the request filters, by translating them to a `WHERE` clause.

### Attention

Only XPath filters are taken into account for index optimization. If the request includes non-optimizable filters, table rows will be fetched from the database, then filtered in Java memory by EBX®, until the requested page size is reached. This is not as efficient as filtering on the database side (especially regarding I/O).

Information on the transmitted SQL request is logged to the category *persistence*. See [Configuring the EBX® logs](#) [p 354].

## Indexing

In order to improve the speed of operations on tables, indexes may be declared on a table at the data model level. This will trigger the creation of an index of the corresponding table in the database.

When designing an index aimed at improving the performance of a given request, the same rules apply as for traditional database index design.

## Setting a fetch size

In order to improve performance, a fetch size should be set according to the expected size of the result of the request on a table. If no fetch size is set, the default value will be used.

- In semantic mode, the default value is 2000.

- In mapped mode, the default value is assigned by the JDBC driver: 10 for Oracle and 0 for PostgreSQL.

### Attention

On PostgreSQL, the default value of 0 instructs the JDBC driver to fetch the whole result set at once, which could lead to an `OutOfMemoryError` when retrieving large amounts of data. On the other hand, using `fetchSize` on PostgreSQL will invalidate server-side cursors at the end of the transaction. If, in the same thread, you first fetch a result set with a `fetchsize`, then execute a procedure that commits the transaction, then, accessing the next result will raise an exception.

### See also

`Request.setFetchSize`<sup>API</sup>

`RequestResult`<sup>API</sup>

---

# Administration Guide

---

## CHAPTER 52

# Administration overview

The Administration section in TIBCO EBX® is the main point of entry for all administration tasks. In this overview are listed all the topics that an administrator needs to master. Click on your topic of interest in order to access the corresponding chapter or paragraph in the documentation.

This chapter contains the following topics:

1. [Repository management](#)
2. [Disk space management](#)
3. [Data model](#)
4. [Perspectives](#)
5. [Administrative delegation](#)

## 52.1 Repository management

For storage optimization, it is recommended to maintain a repository (persistence RDBMS) to the necessary minimum. To this end, it is recommended to regularly perform a purge of snapshots and obsolete dataspaces and to consider using a backup file system.

See also [Cleaning up dataspaces, snapshots, and history](#) [p 379] and [Deleting dataspaces, snapshots, and history](#) [p 380].

It is also possible to archive files of the file system type in order to reduce the storage costs, see [EBX® monitoring](#) [p 378].

Administration tasks can be scheduled by means of the task scheduler, using built-in tasks, see [Task scheduler](#) [p 415].

### **Object cache**

EBX® maintains an object cache in memory. The object cache size should be managed on a case by case basis according to specific needs and requirements (pre-load option and pre-validate on the reference dataspaces, points of reference, and monitoring), while continuously monitoring the repository health reports (./ebxLog/monitoring.log).

See [Memory management](#) [p 296].

### **Obsolete contents**

Keeping obsolete contents in the repository can lead to a slow server startup and slow responsiveness of the interface. It is strongly recommended to delete obsolete content.



For example: datasets referring to deleted data models or undeployed add-on modules. See [Deploying and registering TIBCO EBX® add-ons](#) [p 371].

## Workflow

### Cleanup

The workflow history and associated execution data have to be cleaned up on a regular basis.

The workflow history stores information on completed workflows, their respective steps and contexts. This leads to an ever-growing database containing obsolete history and can thus lead to poor performance of the database if not purged periodically. See [Workflow history](#) [p 414] for more information.

### Email configuration

It is required to configure workflow emails beforehand in order to be able to implement workflow email notifications. See [Configuration](#) [p 412] for more information.

## 52.2 Disk space management

### Purge of logs

The log file size will vary according to the log level (and to the selected severity level) and disk space needs to be accordingly managed.

An automatic purge is provided with EBX®, allowing to define how many days should log files be stored. After the defined period, log files are deleted.

Any customized management of the purge of logs (backup, archiving, etc.) is the user's responsibility.

```
#####
## Directory of log files 'ebxFile:'
## This property is used by special appender prefixed
## by 'ebxFile:' (see log section below)
#####
ebx.logs.directory=${ebx.home}/ebxLog

#####
# Daily rollover threshold of log files 'ebxFile:'
# Specifies the maximum number of backup files for daily rollover of 'ebxFile:' appenders.
# When set to a negative value, backup log files are never purged.
# Default value is -1.
#####
ebx.log4j.appender.ebxFile.backup.Threshold=-1
```

### Audit trail

EBX® is provided with a default audit trail manager. Any customized management (including purge, backups, etc.) is the user's responsibility.

If the audit trail is unwanted, it is possible to fully deactivate it. See [Activating the XML audit trail](#) [p 353] and [Audit trail](#) [p 421] for more information.

## 52.3 Data model

### ***Publication management***

The management of publications of [embedded data models](#) [p 83]. See [Data model administration](#) [p 405] for more information on the management of these publications and the administration tasks that can be performed (delete, import and export).

### ***Refresh data models***

It is possible to update the data models that are using XML Schema documents not managed by EBX®. See [Data model refresh tool](#) [p 473] for more information.

## 52.4 Perspectives

EBX® offers extensive UI customization options. Simplified interfaces ([Recommended perspectives](#)) [p 397] dedicated to each profile accessing the system can be parameterized by the administrator. According to the profile of the user logging in, the interface will offer more or less options and menus. This allows for a streamlined work environment.

See [Advanced perspective](#) [p 386] for more information.

## 52.5 Administrative delegation

EBX® is provided with the built-in administrator profile by default. An administrator can delegate administrative rights to a non-administrator user, either for specific actions or for all activities.

The administrative delegation is defined under 'Administration' in the [global permissions](#) [p 385] profile.

Access to the administration section can be granted to specific profiles via the global permissions in order to delegate access rights on corresponding administration datasets.

If all necessary administrative rights have been delegated to non-administrator users, it becomes possible to disable the built-in 'Administrator' role.

See also [Configuring the user and roles directory](#) [p 352]

---

# Installation & configuration

---

## CHAPTER 53

# Supported environments

This chapter contains the following topics:

1. [Browsing environment](#)
2. [Supported application servers](#)
3. [Supported databases](#)

## 53.1 Browsing environment

### ***Supported web browsers***

The TIBCO EBX® web interface supports the following browsers:

|                                             |                                                                                                                                                                                                         |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Microsoft Edge</b>                       | Minimum supported version is 44<br>Compatibility mode is not supported.                                                                                                                                 |
| <b>Mozilla Firefox ESR 68 (see details)</b> | As Mozilla Firefox is updated frequently, Cloud Software Group, Inc. only fully supports version ESR 68. See <a href="#">Mozilla Firefox ESR</a> for more details.                                      |
| <b>Google Chrome</b>                        | As Google Chrome is updated frequently and it is not possible to deactivate automatic updates, Cloud Software Group, Inc. only tests and makes the best effort to support the latest version available. |

### ***Screen resolution***

The minimum screen resolution for EBX® is 1024x768.

### ***Refreshing pages***

Browser page refresh is not supported by EBX®. When a page refresh is performed, the last user action is re-executed, and therefore could cause issues. It is thus imperative to use the action buttons and links offered by EBX® instead of refreshing the page.

## 'Previous' and 'Next' buttons

The 'previous' and 'next' buttons of the browser are not supported by EBX®. When navigating through page history, an obsolete user action is re-executed, and therefore could cause issues. It is thus imperative to use the action buttons and links offered by EBX® rather than the browser buttons.

## Zoom troubleshooting

Zooming in or out may cause some minor display issues (for example extra scrollbar or misalignment). Those issues can be fixed by refreshing the screen using the provided navigation links.

## Browser configuration

The following features must be activated in the browser configuration, for the user interface to work properly:

- JavaScript
- Ajax
- Pop-ups

### Attention

Avoid using any browser extensions or plug-ins, as they could interfere with the proper functioning of EBX®.

## 53.2 Supported application servers

EBX® supports the following configurations:

- Java Runtime Environment: JRE 17 LTS.

### Note

JRE 8 and 11 are no longer supported. We recommend upgrading your JRE to Java 17 LTS version to take advantage of its performance improvements.

### Note

The [Unicode CLDR](#) has changed with JRE versions and impacts locale-sensitive formats such as integer or decimal representations. Configure the search order of locale sensitive services using the `java.locale.providers` system property. You can use the value `COMPAT` to ensure compatibility with JRE 8. See [java.util.spi.LocaleServiceProvider](#) for more details.

- Any Servlet/JSP container that complies with Servlet 3.1 up to, but not including, 5.0. For example Tomcat 9.0 up to, but not including, 10.0. Any Java Application Server that supports Java SE 17 and does not use Jakarta EE 9 or greater. For example, WebSphere Liberty 23 or higher for Java EE 8, WebLogic Application Server 14c 14.1.2 (2024) or higher, JBoss EAP 7.4 Update 8+ or higher. See [Java EE deployment overview](#) [p 322].
- The application server must use UTF-8 encoding for HTTP query strings from EBX®. This can be set at the application server level.

For example, on Tomcat, you can set the server to always use the UTF-8 encoding, by setting `URIEncoding` to 'UTF-8' on the `<Connector>` in the `server.xml` configuration file. Alternatively, you can instruct the server to use the encoding of the request body by setting the parameter `useBodyEncodingForURI` to 'true' in `server.xml`.

### Attention

- Limitations apply regarding clustering and hot deployment/undeployment:

Clustering: EBX® does not include a cache synchronization mechanism, thus it cannot be deployed into a cluster of active instances. See [Technical architecture](#) [p 374] for more information.

Hot deployment/undeployment: EBX® does not support hot deployment/undeployment of web applications registered as EBX® modules, or of EBX® built-in web applications.

## 53.3 Supported databases

The EBX® repository supports the relational database management systems listed below, with the suitable JDBC drivers. It is important to follow the database vendor recommendations and update policies regarding the database itself, as well as the JDBC driver.

|                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Oracle Database 19c, 23c or higher.</b>                                           | <p>The distinction of null values bears certain limitations. On simple <code>xs:string</code> elements, Oracle does not support the distinction between empty strings and null values. See <a href="#">Empty string management</a> [p 529] for more information.</p> <p>The user with which EBX® connects to the database requires the following privileges:</p> <ul style="list-style-type: none"> <li>• CREATE SESSION,</li> <li>• CREATE TABLE,</li> <li>• ALTER SESSION,</li> <li>• CREATE SEQUENCE,</li> <li>• A non-null quota on its default tablespace.</li> </ul>                                                                                                      |
| <b>PostgreSQL 12 or higher.</b>                                                      | <p>When using PostgreSQL as the underlying database, a request fetch size must be set, otherwise the JDBC driver will fetch the whole result set at once. This could lead to an <code>OutOfMemoryError</code> when retrieving large amounts of data.</p> <p>Also, see this <a href="#">limitation</a> [p 266] regarding the evolution of datamodels in mapped modes.</p> <p>See <code>Request.setFetchSize<sup>API</sup></code>.</p> <p>The user with which EBX® connects to the database needs the <code>CONNECT</code> privilege on the database hosting the EBX® repository. Other than this, the default privileges on the public schema of this database are suitable.</p> |
| <b>Amazon Aurora PostgreSQL (compatible with PostgreSQL 12.13 or higher).</b>        | The comments in the above section for PostgreSQL apply.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Google Cloud SQL for PostgreSQL (compatible with PostgreSQL 12.16 or higher).</b> | The comments in the above section for PostgreSQL apply.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Microsoft SQL Server 2014 or higher.</b>                                          | <p>When used with Microsoft SQL Server, EBX® uses the default database collation to compare and sort strings stored in the database. This applies to strings used in the data model definition, as well as data stored in relational and history tables. The default database collation can be specified when the database is created. Otherwise, the collation of</p>                                                                                                                                                                                                                                                                                                          |

the database server is used. To avoid naming conflicts or unexpected behaviors, a case- and accent-sensitive collation must be used as the default database collation (the collation name is suffixed by "CS\_AS" or the collation is binary).

The default setting to enforce transaction isolation on SQL Server follows a pessimistic model. Rows are locked to prevent any read/write concurrent accesses. This may cause liveliness issues for mapped tables (history or relational). To avoid such issues, it is recommended to activate [snapshot isolation](#) on your SQL Server database.

The user with which EBX® connects to the database requires the following privileges:

- CONNECT, SELECT and CREATE TABLE on the database hosting the EBX® repository,
- ALTER, CONTROL, UPDATE, INSERT, DELETE on its default schema.

### Microsoft Azure SQL Database

EBX® has been qualified on Microsoft Azure SQL Database v12 (12.00.700), and is regularly tested to verify compatibility with the current version of the Azure database service.

When used with Microsoft Azure SQL, EBX® uses the default database collation to compare and sort strings stored in the database. This applies to strings used in the data model definition, as well as data stored in relational and history tables. The default database collation can be specified when the database is created. Otherwise, the database engine server collation is used. To avoid naming conflicts or unexpected behaviors, a case- and accent-sensitive collation must be used as the default database collation (the collation name is suffixed by "CS\_AS" or the collation is binary).

The user with which EBX® connects to the database requires the following privileges:

- CONNECT, SELECT and CREATE TABLE on the database hosting the EBX® repository,
- ALTER, CONTROL, UPDATE, INSERT, DELETE on its default schema.

### H2 2.2.224 or higher.

H2 is not supported for production environments.

The default H2 database settings do not allow consistent reads when records are modified. Relational tables are locked following a pessimistic model. To prevent concurrency issues, it is possible to activate the [MVCC feature](#). Note, however, that the H2 documentation states this feature is not yet fully tested.



For other relational databases, please contact the Support team at <https://support.tibco.com>.

**Attention**

In order to guarantee the integrity of the EBX® repository, it is strictly forbidden to perform direct modifications to the database (for example, using direct SQL writes), except in the specific use cases described in the section [SQL access to data in relational mode](#) [p 246].

**See also**

[Repository administration](#) [p 374]

[Data source of the EBX® repository](#) [p 320]

[Configuring the EBX® repository](#) [p 351]



## CHAPTER 54

# Java EE deployment

This chapter contains the following topics:

1. [Introduction](#)
2. [Software components](#)
3. [Embedded third-party libraries](#)
4. [Required third-party libraries](#)
5. [Web applications](#)
6. [Deployment details](#)
7. [Installation notes](#)

## 54.1 Introduction

This chapter details deployment specifications for TIBCO EBX® on a Java application server. For specific information regarding supported application servers and inherent limitations, see [Supported environments](#). [p 308]

## 54.2 Software components

EBX® uses the following components:

- Library `ebx.jar`
- [Embedded](#) [p 316] and [required](#) [p 316] third-party Java libraries
- [EBX® built-in web applications](#) [p 319] and optional [custom web applications](#) [p 319]
- [EBX® main configuration file](#) [p 349]
- [EBX® repository](#) [p 374]
- [Default user and roles directory](#) [p 401], integrated within the EBX® repository, or a third-party system (LDAP, RDBMS) for the user authentication

See also [Supported environments](#) [p 308]

## 54.3 Embedded third-party libraries

To increase EBX® independence and interoperability, it embeds its own third-party libraries. Even if some of them have been modified, preventing conflicts, others must remain unchanged since they are official Java APIs.

The ones that can produce conflicts are:

- Apache Geronimo JSON
- Javax Activation
- Javax Annotations
- Javax JSON Bind
- Javax SAAJ API
- Javax WS RS
- Javax XML Bind

For more information regarding the versions or the details of the Third-Party Library, please refer to the: `TIB_ebx_5.9.26_license.pdf`.

Since those libraries are already integrated, custom web applications should not include them anew, otherwise linkage errors can occur. Furthermore, they should not be deployed aside from the `ebx.jar` library for the same reasons.

## 54.4 Required third-party libraries

EBX® requires several third-party Java libraries. These libraries must be deployed and be accessible from the class-loader of `ebx.jar`. Depending on the application server and the Java runtime environment being used, these libraries may already be present or may need to be added manually.

### ***Database drivers***

The EBX® repository requires a database. Generally, the required driver is configured along with a data source, if one is used. Depending on the database defined in the main configuration file, one of

the following drivers is required. Keep in mind that, whichever database you use, the version of the JDBC client driver must be equal to or higher than the version of the database server.

|                        |                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>H2</b>              | <p>Version 2.2.224 validated. Note that H2 is not supported in production environments.</p> <p><a href="https://www.h2database.com/">https://www.h2database.com/</a></p> <p>If the repository has been created using H2 version 1.x, <a href="#">this process</a> must be applied, to migrate to the H2 v2 format.</p>                                                           |
| <b>Oracle JDBC</b>     | <p>Oracle database 23ai is validated on their latest patch set update.</p> <p>Determine the driver that should be used according to the database server version and the Java runtime environment version. Download the ojdbc11.jar certified library with JDK 11.</p> <p><a href="#">Oracle database JDBC drivers download.</a></p>                                              |
| <b>SQL Server JDBC</b> | <p>SQL Server 2014 and greater, with all corrective and maintenance patches applied, are validated.</p> <p>Remember to use an up-to-date JDBC driver, as some difficulties have been encountered with older versions.</p> <p>Include the mssql-jdbc-12.6.1.jre11.jar library with JDK 11.</p> <p><a href="#">Download Microsoft JDBC Driver 12.6.1 for SQL Server (zip).</a></p> |
| <b>PostgreSQL</b>      | <p>PostgreSQL Java 8 validated</p> <p>Include the latest JDBC driver version 4.2 released for your database server and Java runtime environment.</p> <p><a href="#">PostgreSQL JDBC drivers download.</a></p>                                                                                                                                                                    |

#### See also

[Data source of the EBX® repository](#) [p 320]

[Configuring the EBX® repository](#) [p 351]

## SMTP and emails

The library for JavaMail 1.5.6 email management is required.

The following libraries are used by email features in EBX®. See [Activating and configuring SMTP and emails](#) [p 357] for details on the configuration.

- javax.mail.jar, version 1.5.6, from August 10, 2016
- smtp.jar, version 1.5.6, from August 10, 2016
- pop3.jar, version 1.5.6, from August 10, 2016

See also [JavaMail](#)

## Secure Socket Layer (SSL)

These libraries are required if your web applications use SSL features.

- jsse.jar: <https://www.oracle.com/java/technologies/jsse-v103-for-cdc-v102.html>
- ibmjse.jar: <https://www.ibm.com/developerworks/java/jdk/security/>

See also [TIBCO EBX® main configuration file](#) [p 349]

## Java Message Service (JMS)

When using JMS, version 1.1 or higher is required.

Depending on whether a Java EE application server or a Servlet/Java Server Pages (JSP) implementation is being used, the library required is as follows:

- For an application server based on Java EE (Java Platform Enterprise Edition), the required JMS provider library is available by default. See <http://www.oracle.com/technetwork/java/javasee/overview> for more information.
- For a Servlet/Java Server Pages (JSP) implementation using Java SE (Java Platform Standard Edition), for example Apache Tomcat, a JMS provider library such as [Apache ActiveMQ](#) may need to be added. See <http://www.oracle.com/technetwork/java/javase/overview> for more information.

### Note

In EBX®, the supported JMS model is exclusively Point-to-Point (PTP). PTP systems allow working with queues of messages.

See also [TIBCO EBX® main configuration file](#) [p 349]

## 54.5 Web applications

EBX® provides pre-packaged EARs that can be deployed directly if your company has no custom EBX® module web applications to add. If deploying custom web applications as EBX® modules, it is recommended to rebuild an EAR containing the custom modules packaged at the same level as the built-in web applications.

### Attention

Web application deployment on / path context is no more supported. The path context must not be empty nor equals to /. Moreover, web applications deployment on paths of different depth is deprecated. Every web application path context must be set on the same path depth.

For more information, see the note on [repackaging the EBX® EAR](#) [p 323] at the end of this chapter.

## EBX® built-in web applications

EBX® includes the following built-in web applications.

| Web application name | Description                                                                                                                                                                                                                                                                                                                                                                                                                                | Required |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| ebx                  | EBX® entry point, which handles the initialization on start up. See <a href="#">Deployment details</a> [p 319] for more information.                                                                                                                                                                                                                                                                                                       | Yes      |
| ebx-root-1.0         | EBX® root web application. Any application that uses EBX® requires the root web application to be deployed.                                                                                                                                                                                                                                                                                                                                | Yes      |
| ebx-ui               | EBX® user interface web application.                                                                                                                                                                                                                                                                                                                                                                                                       | Yes      |
| ebx-manager          | EBX® user interface web application.                                                                                                                                                                                                                                                                                                                                                                                                       | Yes      |
| ebx-dma              | EBX® data model assistant, which helps with the creation of data models through the user interface.<br><b>Note:</b> The data model assistant requires the ebx-manager user interface web application to be deployed.                                                                                                                                                                                                                       | Yes      |
| ebx-dataservices     | EBX® data services web application. Data services allow external interactions with the EBX® repository using the <a href="#">SOAP operations</a> [p 617] and Web Services Description Language <a href="#">WSDL generation</a> [p 609] standards or using the <a href="#">Built-in RESTful services</a> [p 659].<br><b>Note:</b> The EBX® web service generator requires the deployment of the ebx-manager user interface web application. | Yes      |

## Custom web applications

It is possible to extend and customize the behavior of EBX® by deploying custom web applications which conform to the EBX® module requirements.

See also

[Packaging TIBCO EBX® modules](#) [p 461]

[Declaring modules as undeployed](#) [p 364]

## 54.6 Deployment details

### Introduction

This section describes the various options available to deploy the 'ebx' web application. These options are available in its deployment descriptor (`WEB-INF/web.xml`) and are complemented by the properties defined in the main configuration file.

#### Attention

For JBoss application servers, any unused resources must be removed from the `WEB-INF/web.xml` deployment descriptor.

**See also**

[TIBCO EBX® main configuration file](#) [p 349]

[Supported application servers](#) [p 309]

**User interface and web access**

The web application 'ebx' (packaged as ebx.war) contains the servlet `FrontServlet`, which handles the initialization and serves as the sole user interface entry point for the EBX® web tools.

**Configuring the deployment descriptor for 'FrontServlet'**

In the file `WEB-INF/web.xml` of the web application 'ebx', the following elements must be configured for `FrontServlet`:

|                                                   |                                                                                                                                                                                                                  |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>/web-app/servlet/load-on-startup</code>     | To ensure that <code>FrontServlet</code> initializes upon EBX® start up, the <code>web.xml</code> deployment descriptor must specify the element <code>&lt;load-on-startup&gt;1&lt;/load-on-startup&gt;</code> . |
| <code>/web-app/servlet-mapping/url-pattern</code> | <code>FrontServlet</code> must be mapped to the path <code>'/'</code> .                                                                                                                                          |

**Configuring the application server for 'FrontServlet'**

- `FrontServlet` must be authorized to access other contexts, such as `ServletContext`.

For example, on Tomcat, this configuration is performed using the attribute `crossContext` in the configuration file `server.xml`, as follows:

```
<Context path="/ebx" docBase="(...)" crossContext="true"/>
```

**Data source of the EBX® repository****Note**

If the EBX® main configuration specifies the property `ebx.persistence.url`, then the environment entry below will be ignored by EBX® runtime. This option is only provided for convenience; it is always recommended to use a fully-configurable datasource. See [Configuring the EBX® repository](#) [p 351] for more information on this property.

The JDBC datasource for EBX® is specified in the deployment descriptor `WEB-INF/web.xml` of the 'ebx' web application as follows:

| Reserved resource name           | Default JNDI name                                                                 | Description                                                                           |
|----------------------------------|-----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <code>jdbc/EBX_REPOSITORY</code> | Weblogic: <code>EBX_REPOSITORY</code><br>JBoss: <code>java:/EBX_REPOSITORY</code> | JDBC data source for EBX® Repository.<br>Java type: <code>javax.sql.DataSource</code> |

**See also**

[Configuring the EBX® repository](#) [p 351]



[Rules for the database access and user privileges](#) [p 375]

## Mail sessions

### Note

If the EBX® main configuration does not set `ebx.mail.activate` to 'true', or if it specifies the property `ebx.mail.smtp.host`, then the environment entry below will be ignored by EBX® runtime. See [SMTP](#) [p 357] in the EBX® main configuration properties for more information on these properties.

SMTP and email is declared in the deployment descriptor `WEB-INF/web.xml` of the 'ebx' web application as follows:

| Reserved resource name | Default JNDI name                                                      | Description                                                                                        |
|------------------------|------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| mail/EBX_MAIL_SESSION  | Weblogic:<br>EBX_MAIL_SESSION<br><br>JBoss: java:/<br>EBX_MAIL_SESSION | Java Mail session used to send emails from EBX®.<br><br>Java type: <code>javax.mail.Session</code> |

## JMS connection factory

### Note

If the EBX® main configuration does not activate JMS through the property `ebx.jms.activate`, the environment entry below will be ignored by the EBX® runtime. See [JMS](#) [p 358] in the EBX® main configuration properties for more information on this property.

The JMS connection factory is declared in the deployment descriptor `WEB-INF/web.xml` of the 'ebx' web application as follows:

| Reserved resource name       | Default JNDI name                                                                      | Description                                                                                                                                                                                                        | Required |
|------------------------------|----------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| jms/EBX_JMSConnectionFactory | Weblogic:<br>EBX_JMSConnectionFactory<br><br>JBoss: java:/<br>EBX_JMSConnectionFactory | JMS connection factory used by EBX® to create connections with the JMS provider configured in the operational environment of the application server.<br><br>Java type:<br><code>javax.jms.ConnectionFactory</code> | Yes      |

### Note

For deployment on WildFly, JBoss and WebLogic application servers with JNDI capabilities, you must update `EBX.ear` or `EBXForWebLogic.ear` for additional mappings of all required resource names to JNDI names.

## JMS for data services

To configure data services to use JMS instead of the default HTTP, you must configure the [JMS connection factory](#) [p 321] and the following queues, declared in the `WEB-INF/web.xml` deployment descriptor of the 'ebx' web application. This is the only method for configuring JMS for data services.

When a SOAP request is received, the SOAP response is optionally returned if the header field `JMSReplyTo` is defined. If so, the fields `JMSCorrelationID` and `JMSType` are retained.

See [JMS](#) [p 358] for more information on the associated EBX® main configuration properties.

### Note

If the EBX® main configuration does not activate JMS through the property `ebx.jms.activate`, then the environment entries below will be ignored by EBX® runtime. See [JMS](#) [p 358] in the EBX® main configuration properties for more information on this property.

| Reserved resource name            | Default JNDI name                                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Required |
|-----------------------------------|-------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| <code>jms/EBX_QueueIn</code>      | Weblogic: <code>EBX_QueueIn</code><br>JBoss: <code>java:/jms/EBX_QueueIn</code>           | JMS queue for incoming SOAP requests sent to EBX® by other applications.<br>Java type: <code>javax.jms.Queue</code>                                                                                                                                                                                                                                                                                                                                                                            | No       |
| <code>jms/EBX_QueueFailure</code> | Weblogic: <code>EBX_QueueFailure</code><br>JBoss: <code>java:/jms/EBX_QueueFailure</code> | JMS queue for failures. It contains incoming SOAP requests for which an error has occurred. This allows replaying these messages if necessary.<br>Java type: <code>javax.jms.Queue</code><br><b>Note:</b> For this property to be read, the main configuration must also activate the queue for failures through the property <code>ebx.jms.activate.queueFailure</code> . See <a href="#">JMS</a> [p 358] in the EBX® main configuration properties for more information on these properties. | No       |

## JAR files scanner

To speed up the web applications server startup, the JAR files scanner configuration should be modified to exclude, at least, the `ebx.jar` and `ebx-addons.jar` libraries.

For example, on Tomcat, this should be performed in the `tomcat.util.scan.DefaultJarScanner.jarToSkip` property from the `catalina.properties` file.

## 54.7 Installation notes

EBX® can be deployed on any Java EE 7 to Jakarta EE 8 application server that supports Servlet 3.1, up to but not including version 5.0. The following documentation on Java EE deployment and installation notes are available:

- [Installation note for Tomcat 9.x](#) [p 331]
- [Installation note for JBoss EAP 7.4.x](#) [p 325]
- [Installation note for WebSphere Application Server Liberty 23.x](#) [p 337]

- [Installation note for WebLogic 14c](#) [p 343]

### Attention

- The EBX® installation notes on Java EE 7 to Jakarta EE 8 application servers do not replace the native documentation for each application server.
- These are *not* general installation recommendations, as the installation process is determined by architectural decisions, such as the technical environment, application mutualization, delivery process, and organizational decisions.
- In these examples, no additional EBX® modules are deployed. To deploy additional modules, the best practice is to rebuild an EAR with the module as a web application at the same level as the other EBX® modules. The web application must declare its class path dependency as specified by the Java™ 2 Platform Enterprise Edition Specification, v1.4:

#### **J2EE.8.2 Optional Package Support**

(...)

*A JAR format file (such as a JAR file, WAR file, or RAR file) can reference a JAR file by naming the referenced JAR file in a Class-Path header in the Manifest file of the referencing JAR file. The referenced JAR file is named using a URL relative to the URL of the referencing JAR file. The Manifest file is named META-INF/MANIFEST.MF in the JAR file. The Class-Path entry in the Manifest file is of the form:*

*Class-Path: list-of-jar-files-separated-by-spaces*

In an "industrialized" process, it is strongly recommended to develop a script that automatically builds the EAR, with the custom EBX® modules, the EBX® web applications, as well as all the required shared libraries.

- In order to avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar` or other libraries in the class-loading system.
- In case of deployment on Oracle WebLogic server, please refer to the [Module structure](#) [p 461] section.



## CHAPTER 55

# Installation note for JBoss EAP 7.4.x

This chapter contains the following topics:

1. [Overview](#)
2. [Requirements](#)
3. [JBoss Application Server installation](#)
4. [EBX® home directory configuration](#)
5. [JBoss Application Server and Java Virtual Machine configuration](#)
6. [JNDI entries configuration](#)
7. [Data source and JDBC provider configuration](#)
8. [EBX.ear application update](#)
9. [EBX.ear application deployment](#)
10. [EBX® application start](#)

## 55.1 Overview

### Attention

- This chapter describes a *quick installation example* of TIBCO EBX® on JBoss Application Server.
- It does not replace the [documentation](#) of this application server.
- They are *not* general installation recommendations, as the installation process is determined by architectural decisions, such as the technical environment, application mutualization, delivery process, and organizational decisions.
- The complete description of the components needed by EBX® is given in chapter [Java EE deployment](#) [p 315].
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar` or other libraries in the class-loading system.
- Refer to the [Security Best Practices](#) [p 454] for information on navigating secure deployments.

## 55.2 Requirements

- Java SE 17 LTS
- JBoss Application Server EAP 7.4 ([update 8 and above](#))
- Database and JDBC driver
- EBX® CD
- No CDI features in EBX®'s additional modules (since CDI will be automatically disable)

See also [Supported environments](#) [p 308]

## 55.3 JBoss Application Server installation

*This quick installation example is performed for an unix operating system.*

1. Download JBoss EAP 7.4.0 Installer jar version from:  
<https://developers.redhat.com/products/eap/download/>
2. Run the Installer using `java -jar` command line.  
For further installation details, refer to the [documentation](#).
3. Perform a standard installation:
  1. Select the language and click 'OK',
  2. Accept the License and click 'Next',
  3. Choose the installation path and click 'Next',
  4. Keep the 'Component Selection' as it is and click 'Next',
  5. Enter 'Admin username', 'Admin password' and click 'Next',
  6. On 'Installation Overview' click 'Next',
  7. On 'Component Installation' click 'Next',
  8. On 'Configure Runtime Environment' leave the selection as it is and click 'Next',
  9. When 'Processing finished' appear, click 'Next',
  10. Uncheck 'Create shortcuts in the start menu' and click 'Next',
  11. Generate 'installation script and properties file' in the JBoss EAP 7.4 installation root directory,
  12. Click 'done'.
4. Download and apply the last update for JBoss EAP 7.4.0, this link requires an authentication:  
<https://access.redhat.com/jbossnetwork/restricted/listSoftware.html?downloadType=patches&product=appplatform&version=7.4>

## 55.4 EBX® home directory configuration

1. Create the `EBX_HOME` directory, for example `/opt/ebx/home`.

2. Copy from the *EBX® CD*, the `ebx.software/files/ebx.properties` file to *EBX\_HOME*. In our example, we will have the following file:  
`/opt/ebx/home/ebx.properties.`
3. If needed, edit the `ebx.properties` file to override the default database. By default the standalone H2 database is defined. The property key `ebx.persistence.factory` must be uncommented for other supported databases and the `h2.standalone` one must be commented.

## 55.5 JBoss Application Server and Java Virtual Machine configuration

1. Open the `standalone.conf` configuration file, placed in `<JBoss_HOME>/bin` (or `jboss-eap.conf` file placed in `<JBoss_HOME>/bin/init.d` for running the server as a service).
2. Add 'ebx.properties' and 'ebx.home' properties to the 'JAVA\_OPTS' environment variable respectively set with `ebx.properties` file's path and *EBX\_HOME* directory's path.

## 55.6 JNDI entries configuration

1. Open the `standalone-full.xml` file placed in `<JBoss_HOME>/standalone/configuration`.
2. Add, at least, the following lines to the `server` tag in `messaging-activemq` subsystem:

```
<connection-factory
  name="jms/EBX_JMSConnectionFactory"
  entries="java:/EBX_JMSConnectionFactory"
  connectors="To Be Defined"/>
<jms-queue
  name="jms/EBX_D3ReplyQueue"
  entries="java:/jms/EBX_D3ReplyQueue"
  durable="true"/>
<jms-queue
  name="jms/EBX_QueueIn"
  entries="java:/jms/EBX_QueueIn"
  durable="true"/>
<jms-queue
  name="jms/EBX_QueueFailure"
  entries="java:/jms/EBX_QueueFailure"
  durable="true"/>
<jms-queue
  name="jms/EBX_D3MasterQueue"
  entries="java:/jms/EBX_D3MasterQueue"
  durable="true"/>
<jms-queue
  name="jms/EBX_D3ArchiveQueue"
  entries="java:/jms/EBX_D3ArchiveQueue"
  durable="true"/>
<jms-queue
  name="jms/EBX_D3CommunicationQueue"
  entries="java:/jms/EBX_D3CommunicationQueue"
  durable="true"/>
```

*Caution:* the `connectors` attribute value, from the `connection-factory` element, has to be defined. Since the kind of connectors is strongly reliant on the environment infrastructure, a default configuration can not be provided.

See [configuring messaging](#) for more information.

3. Add, at least, the following line to `mail` subsystem:

```
<mail-session name="mail" debug="false" jndi-name="java:/EBX_MAIL_SESSION"/>
```

## 55.7 Data source and JDBC provider configuration

1. After the launch of the JBoss Server, run the management CLI without the use of '--connect' or '-c' argument.
2. Use the 'module add' management CLI command to add the new core module. Sample for PostgreSQL configuration:

```
module add \
--name=org.postgresql \
--resources=<PATH_TO_JDBC_JAR> \
--dependencies=javaee.api, sun.jdk, ibm.jdk, javax.api, javax.transaction.api
```

3. Use the 'connect' management CLI command to connect to the running instance.
4. Register the JDBC driver. When running in a managed domain, ensure to precede the command with '/profile=<PROFILE\_NAME>'. Sample for PostgreSQL configuration:

```
/subsystem=\
datasources/jdbc-driver=\
  postgresql:add(\
    driver-name=postgresql, \
    driver-module-name=org.postgresql, \
    driver-xa-datasource-class-name=org.postgresql.xa.PGXADatasource\
  )
```

5. Define the datasource using the 'data-source add' command, specifying the appropriate argument values. Sample for PostgreSQL configuration:

```
data-source add \
--name=jdbc/EBX_REPOSITORY \
--jndi-name=java:/EBX_REPOSITORY \
--driver-name=postgresql \
--connection-url=jdbc:postgresql://<SERVER_NAME>:<PORT>/<DATABASE_NAME> \
--user-name=<PERSISTENCE_USER> \
--password=<PERSISTENCE_PASSWORD>
```

## 55.8 EBX.ear application update

1. Copy from the *EBX® CD*, the `ebx.software/webapps/ear-packaging/EBX.ear` file to your working directory.
2. Uncompress the ear archive to add the application's specific required third-party libraries and additional web modules.

Mail: see [SMTP and emails](#) [p 317] for more information.

SSL: see [Secure Socket Layer \(SSL\)](#) [p 318] for more information.

JMS: see [Java Message Service \(JMS\)](#) [p 318] for more information.

3. Update the `/META-INF/application.xml` and `/META-INF/jboss-deployment-structure.xml` files according to the added additional web modules.
4. Compress anew the ear archive.

## 55.9 EBX.ear application deployment

1. Copy `EBX.ear` into the `JBoss_HOME/standalone/deployments` directory.



## 55.10 EBX® application start

1. After the launch of the JBoss Application Server, with the `<JBoss_HOME>/bin/standalone.sh -c standalone-full.xml` command line or through the service command, run the EBX® web application by entering the following URL in the browser: <http://localhost:8080/ebx/>.
2. At first launch, [EBX® Wizard](#) [p 369] helps to configure the default properties of the initial repository.



## CHAPTER 56

---

# Installation note for Tomcat 9.x

This chapter contains the following topics:

1. [Overview](#)
2. [Requirements](#)
3. [Tomcat Application Server installation](#)
4. [EBX® home directory configuration](#)
5. [Tomcat Application Server and Java Virtual Machine configuration](#)
6. [EBX® and third-party libraries deployment](#)
7. [EBX® web applications deployment](#)
8. [EBX® application start](#)

## 56.1 Overview

### Attention

- This chapter describes a *quick installation example* of TIBCO EBX® on Tomcat Application Server.
- It does not replace the [documentation](#) of this application server.
- They are *not* general installation recommendations, as the installation process is determined by architectural decisions, such as the technical environment, application mutualization, delivery process, and organizational decisions.
- Tomcat 10.x is not supported.
- The complete description of the components needed by EBX® is given in chapter [Java EE deployment](#) [p 315].
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar` or other libraries in the class-loading system.
- The description below uses the variable name `$CATALINA_HOME` to refer to the Tomcat installation directory, and from which most relative paths are resolved. However, if the `$CATALINA_BASE` directory has been set for a multiple instances configuration, it should be used for each of these references.
- Refer to the [Security Best Practices](#) [p 454] for information on navigating secure deployments.

## 56.2 Requirements

- Java SE 17 LTS
- Apache Tomcat 9.x
- Database and JDBC driver
- EBX® CD

See also [Supported environments](#) [p 308]

## 56.3 Tomcat Application Server installation

1. Download Tomcat 9.x core binary distributions from:  
<https://tomcat.apache.org/download-90.cgi>
2. Run the installer or extract the archive and perform a standard installation with default options

## 56.4 EBX® home directory configuration

1. Create `EBX_HOME` directory, for example `C:\EBX\home`, or `/home/ebx`
2. Copy from *EBX® CD* the `ebx.software/files/ebx.properties` file to `EBX_HOME`. In our example, we will have the following file:

C:\EBX\home\ebx.properties, or /home/ebx/ebx.properties

3. If needed, edit the ebx.properties file to override the default database. By default the standalone H2 database is defined. The property key ebx.persistence.factory must be uncommented for other supported databases and the h2.standalone one must be commented.

## 56.5 Tomcat Application Server and Java Virtual Machine configuration

1. Modify \$CATALINA\_HOME/conf/server.xml (or \$CATALINA\_BASE/conf/server.xml) file.

Add the attribute encodedSolidusHandling="passthrough" to the Connector element.

Add the following line to the <Host> element:

```
<Context path="/ebx" crossContext="true" docBase="ebx.war"/>
```

In our example, we will have:

```
<Host name=...>
```

```
... ..
```

```
<Context path="/ebx" crossContext="true" docBase="ebx.war"/>
```

```
... ..
```

```
</Host>
```

2. Modify the \$CATALINA\_HOME/conf/catalina.properties (or \$CATALINA\_BASE/conf/catalina.properties) file by adding the following lines to the tomcat.util.scan.DefaultJarScanner.jarsToSkip property:

```
ebx.jar,\
```

```
ebx-addons.jar,\
```

3. Configure the Java Virtual Machine properties

- For Windows' Command Prompt launch

Set the environment variables by creating a setenv.bat file either into \$CATALINA\_HOME\bin or \$CATALINA\_BASE\bin. This file will hold, at least, the following lines:

```
set EBX_HOME=<path_to_the_directory_ebx_home>
set EBX_OPTS=-Debx.home="%EBX_HOME%" -Debx.properties="%EBX_HOME%\ebx.properties"
set JAVA_OPTS=%EBX_OPTS% -Dorg.apache.catalina.connector.CoyoteAdapter.ALLOW_BACKSLASH=true %JAVA_OPTS%
```

- For Windows users that have installed Tomcat as a service

Set Java options through the Tomcat service manager GUI (Java tab).

Be sure to set options on separate lines in the *Java Options* field of the GUI:

```
-Debx.home="<path_to_the_directory_ebx_home>"
-Debx.properties="<path_to_the_directory_ebx_home>\ebx.properties"
```

Update the service using the //US// parameter to set the proper classpath value.

```
C:\> tomcat9 //US//Tomcat9 --Classpath="%CLASSPATH%"
```

- For Unix shell launch

Set the environment variables by creating a setenv.sh file either into \$CATALINA\_HOME/bin or \$CATALINA\_BASE/bin. This file will hold, at least, the following lines:

```
EBX_HOME="<path_to_the_directory_ebx_home>"
EBX_OPTS="-Debx.home=${EBX_HOME} -Debx.properties=${EBX_HOME}/ebx.properties"
```

```
export JAVA_OPTS="${EBX_OPTS} -Dorg.apache.catalina.connector.CoyoteAdapter.ALLOW_BACKSLASH=true
${JAVA_OPTS}"
```

#### Note

The Tomcat option `-Dorg.apache.catalina.connector.CoyoteAdapter.ALLOW_BACKSLASH=true` avoid blocking URLs containing encoded backslashes.

**Caution:** Accounts used to launch EBX® must have create/update/delete rights on *EBX\_HOME* directory.

#### Note

`<path_to_the_directory_ebx_home>` is the directory where we copied `ebx.properties`. In our example, it is `C:\EBX\home`, or `/home/ebx`.

## 56.6 EBX® and third-party libraries deployment

1. Copy third-party libraries from the *EBX® CD* to `$CATALINA_HOME/lib/` (or `$CATALINA_BASE/lib/`) directory. In our example, we will have:

`$CATALINA_HOME/lib/javax.mail-1.5.6.jar` coming from `ebx.software/lib/lib-mail` directory.

`$CATALINA_HOME/lib/h2-2.2.224.jar` (default persistence factory) coming from `ebx.software/lib/lib-h2` directory.

The exact description of these components is given in chapter [Software components](#) [p 315]. Obviously, if those components are already deployed on the class-loading system, they do not have to be duplicated.

2. Copy from *EBX® CD* the `ebx.software/lib/ebx.jar` file to `$CATALINA_HOME/lib/` (or `$CATALINA_BASE/lib/`) directory. In our example, we will have:

`$CATALINA_HOME/lib/ebx.jar`

## 56.7 EBX® web applications deployment

1. Copy from the *EBX® CD* the war files in `ebx.software/webapps/wars-packaging` to the `$CATALINA_HOME/webapps/` (or `$CATALINA_BASE/webapps/`) directory. In our example, we will have:

`$CATALINA_HOME/webapps/ebx.war`: Initialization servlet for EBX® applications

`$CATALINA_HOME/webapps/ebx-root-1.0.war`: Provides a common default module for data models

`$CATALINA_HOME/webapps/ebx-manager.war`: Master Data Management web application

`$CATALINA_HOME/webapps/ebx-dataservices.war`: Data Services web application

`$CATALINA_HOME/webapps/ebx-dma.war`: Data Model Assistant web application

`$CATALINA_HOME/webapps/ebx-ui.war`: User Interface web application

## 56.8 EBX® application start

1. After Tomcat launch, run EBX® web application by entering the following URL in the browser:  
<http://localhost:8080/ebx/>
2. At first launch, [EBX® Wizard](#) [p 369] helps to configure the default properties of the initial repository.





## CHAPTER 57

---

# Installation note for WebSphere Application Server Liberty 23.x

This chapter contains the following topics:

1. [Overview](#)
2. [Requirements](#)
3. [WebSphere Application Server Liberty installation](#)
4. [EBX® home directory configuration](#)
5. [EBX® and third-party libraries deployment](#)
6. [JNDI entries configuration](#)
7. [Data source and JDBC provider configuration](#)
8. [Java Virtual Machine configuration](#)
9. [EBX® application deployment](#)
10. [EBX® application start](#)

## 57.1 Overview

### Attention

- This chapter describes a *quick installation example* of TIBCO EBX® on WebSphere Application Server Liberty.
- It does not replace the [documentation](#) of this application server.
- They are *not* general installation recommendations, as the installation process is determined by architectural decisions, such as the technical environment, application mutualization, delivery process, and organizational decisions.
- The complete description of the components needed by EBX® is given in the chapter [Java EE deployment](#) [p 315].
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar`, or other libraries in the class-loading system.
- The description below uses the variable name `<WLP_HOME>` to refer to the WebSphere Application Server Liberty installation directory, and from which relative paths are resolved.
- Refer to the [Security Best Practices](#) [p 454] for information on navigating secure deployments.

## 57.2 Requirements

- Java SE 17 LTS. We recommended using [IBM Semeru](#), or [Eclipse Adoptium](#) Java distributions.
- WebSphere Application Server Liberty 23.x for Java EE 8
- Database and JDBC driver
- EBX® CD
- No CDI features in EBX®'s additional modules (since CDI will be automatically disabled)

See also [Supported environments](#) [p 308]

## 57.3 WebSphere Application Server Liberty installation

*This quick installation example is performed for a Linux operating system.*

1. Download WebSphere Application Server Liberty 23.x latest version (zip distribution) from:  
<https://www.ibm.com/support/pages/23001-websphere-application-server-liberty-23001>.

The following example uses *Liberty with Java EE 8 Web Profile* zip distribution.

2. Extract the content of this package to a directory on your local machine.
3. Create a new application server (example: 'ebxServer') by running the command:

```
<WLP_HOME>/wlp/bin/server create ebxServer
```

4. Edit `<WLP_HOME>/usr/servers/<SERVER_NAME>/server.xml` file. In our example, `<SERVER_NAME>` is 'ebxServer'.

- Replace the `<feature>webProfile-8.0</feature>` tag with the following lines:

```
<feature>javaMail-1.5</feature>
```

```
<feature>jdbc-4.2</feature>
<feature>jndi-1.0</feature>
<feature>servlet-4.0</feature>
```

- Add the following line into the `<server>` tag:

```
<webContainer deferServletLoad="false" allowQueryParamWithNoEqual="true"/>
```

#### Note

The `deferServletLoad` attribute must be set to `false` since EBX® requires servlets to be initialized on application server startup. By default, WebSphere Application Server Liberty defers servlets loading until a request is received for the associated web application. See [Specifying when servlets are loaded and initialized](#) for more information.

#### Note

The `allowQueryParamWithNoEqual` attribute must be set to `true` since, according to the Servlet API specification, a value of a request parameter must be returned as a `String`, or `null` if the parameter does not exist. Thus, EBX® expects a non `null` value for an existing request parameter.

## 57.4 EBX® home directory configuration

1. Create the *EBX\_HOME* directory, for example `/opt/ebx/home`.
2. Copy from the *EBX® CD*, the `ebx.software/files/ebx.properties` file to *EBX\_HOME*. In our example, we will have the following file:  
`/opt/ebx/home/ebx.properties`.
3. If needed, edit the `ebx.properties` file to override the default database. By default the standalone H2 database is defined. The property key `ebx.persistence.factory` must be uncommented for other supported databases and the `h2.standalone` one must be commented.

## 57.5 EBX® and third-party libraries deployment

1. Create the *EBX\_LIB* directory. For example, `/opt/ebx/home/lib`.
2. Copy third-party libraries from the *EBX® CD*, or from other sources, to the *EBX\_LIB* directory. In our example and for a PostgreSQL database, we will get:  
`postgresql-X.X.X-driver.jar`, coming from another source than the *EBX® CD*.

The complete description of these components is given in the chapter [Java EE deployment](#) [p 315]. If those components are already deployed on the class-loading system, they do not have to be duplicated (ex: `javax.mail 1.5.x` is already present since `javaMail-1.5` feature has been activated for this server).

## 57.6 JNDI entries configuration

1. For *JMS* configuration on WebSphere Application Server Liberty, see [Configuring JMS connection factories](#).

2. For *JavaMail* configuration on WebSphere Application Server Liberty, see [Administering JavaMail on Liberty](#).

## 57.7 Data source and JDBC provider configuration

1. To create the JDBC driver library, edit the `<WLP_HOME>/usr/servers/<SERVER_NAME>/server.xml` file by adding the following lines into the `<server>` tag:

```
<library id="jdbcDriver">
  <file name="{EBX_LIB}/To be completed"/>
</library>
```

*Caution:* the name attribute value, from the `file` tag, has to be completed. In our example, it should be `{EBX_LIB}/postgresql-X.X.X-driver.jar`.

2. To create the JDBC data source, edit the same file and same XML tag by adding the following lines:

```
<dataSource id="ebxRepository" jndiName="jdbc/EBX_REPOSITORY">
  <jdbcDriver libraryRef="jdbcDriver"/>
  <properties.toBeCompleted/>
</dataSource>
```

*Caution:* the `properties` tag, has to be completed. Since the kind of tag is strongly reliant on the JDBC driver used, a default configuration can not be provided. For a PostgreSQL driver, it may look like:

```
<properties.postgresql databaseName="EBXDB" serverName="localhost" portNumber="5432" user="ebx"
password="ebx"/>
```

### Note

To hide the password for a JDBC connection in the `server.xml` file of WebSphere Application Server Liberty, refer to the documentation [Configuring authentication aliases for Liberty](#).

## 57.8 Java Virtual Machine configuration

1. Set server's environment variables by adding the following lines into `<WLP_HOME>/usr/servers/<SERVER_NAME>/server.env` file:

```
EBX_HOME="<path_to_the_directory_ebx_home>"
EBX_LIB="<path_to_the_directory_ebx_lib>"
JAVA_HOME="<path_to_the_java_home>"
```

2. Set Java options by creating a new `jvm.options` file under the `<WLP_HOME>/usr/servers/<SERVER_NAME>/` directory. This file will hold, at least, the following lines:

```
-Debx.home="{EBX_HOME}"
-Debx.properties="{EBX_HOME}/ebx.properties"
```

## 57.9 EBX® application deployment

1. Create the `<EBX_HOME>/ear/` directory.
2. Copy from the *EBX® CD*, the `ebx.software/webapps/ear-packaging/EBX.ear` to the `<EBX_HOME>/ear/` directory. In our example, we will get:

/opt/ebx/home/ear/EBX.ear

3. To deploy EBX.ear, edit the <WLP\_HOME>/usr/servers/<SERVER\_NAME>/server.xml file by adding the following lines into the <server> tag:

```
<application id="ebxApp" location="${EBX_HOME}/ear/{ebx.ear.name}.ear" name="ebxApp" type="ear"></application>
```

## 57.10 EBX® application start

1. After the launch of the WebSphere Application Server Liberty, with the <WLP\_HOME>/wlp/bin/server start ebxServer command line, run the EBX® web application by entering the following URL in the browser:

<http://localhost:9080/ebx/>

2. At first launch, [EBX® Wizard](#) [p 369] helps to configure the default properties of the initial repository.



## CHAPTER 58

---

# Installation note for WebLogic 14c

This chapter contains the following topics:

1. [Overview](#)
2. [Requirements](#)
3. [WebLogic Application Server installation](#)
4. [EBX® home directory configuration](#)
5. [WebLogic Application Server and Java Virtual Machine configuration](#)
6. [EBX® and third-party libraries deployment](#)
7. [Data source and JDBC provider configuration](#)
8. [EBX® application deployment](#)
9. [EBX® application start](#)

## 58.1 Overview

### Attention

- WebLogic 14c supports JDK 11, but not JDK 17 LTS. Until a WebLogic release supports JDK 17 LTS, EBX® will continue to support WebLogic 14c with JDK 11. Please be aware that, according to [Oracle's Statement of Direction](#), future versions of supported software might not be compatible with EBX®.
- This chapter describes a *quick installation example* of TIBCO EBX® on WebLogic Application Server.
- It does not replace the [documentation](#) of this application server.
- They are *not* general installation recommendations, as the installation process is determined by architectural decisions, such as the technical environment, application mutualization, delivery process, and organizational decisions.
- The complete description of the components needed by EBX® is given in chapter [Java EE deployment](#) [p 315].
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar`, or other libraries in the class-loading system.
- Refer to the [Security Best Practices](#) [p 454] for information on navigating secure deployments.

## 58.2 Requirements

- Certified Oracle Java SE 11 ([update 6 and above](#))
- WebLogic Server 14c
- Database and JDBC driver
- EBX® CD

See also [Supported environments](#) [p 308]

## 58.3 WebLogic Application Server installation

1. Download WebLogic 14c latest version from:  
<https://www.oracle.com/middleware/technologies/fusionmiddleware-downloads.html>
2. Run the Oracle Fusion Middleware Weblogic installation wizard using a certified Oracle JDK and the `java -jar` command line
3. Perform a standard installation with default options and choose the appropriate installation directory
4. Leave the 'Automatically launch the Configuration Wizard' option activated to perform the next steps:
  1. Create Domain: choose 'Create a new domain' and specify the domain home directory, then click 'Next'



2. Templates: keep as default and click 'Next'
3. Administrator Account: enter a domain administrator username and password and click 'Next'
4. Domain Mode and JDK: choose the production mode and your JDK installation home and click 'Next'
5. Advanced configuration: check 'Administration server' and 'Topology'. That way, we create two independent domain nodes: an administration one and an application one.  
Click 'Next'
6. Administration Server: enter your administration node name (for example 'AdminServer') and listen port (by default 7001), then click 'Next'
7. Managed Servers: add the application node name (for example 'EbxServer') and listen port (for example 7003), then click 'Next'
8. Clusters: keep as default and click 'Next'
9. Server Templates: keep as default and click 'Next'
10. Machines: keep as default and click 'Next'
11. Configuration Summary: click 'Create'
12. Configuration Process: click 'Next'
13. End Of Configuration: click 'Finish'

## 58.4 EBX® home directory configuration

1. Create *EBX\_HOME* directory, for example `C:\EBX\home`, or `/home/ebx`
2. Copy from *EBX® CD* the `ebx.software/files/ebx.properties` file to *EBX\_HOME*. In our example, we will have the following file:  
`C:\EBX\home\ebx.properties`, or `/home/ebx/ebx.properties`
3. If needed, edit the `ebx.properties` file to override the default database. By default the standalone H2 database is defined. The property key `ebx.persistence.factory` must be uncommented for other supported databases and the `h2.standalone` one must be commented.

## 58.5 WebLogic Application Server and Java Virtual Machine configuration

1. Configure the launch properties for the *Managed Server* (for example 'EbxServer')  
Edit the `<DOMAIN_HOME>/bin/startManagedWebLogic.sh` script file by adding the following lines:

```
EBX_HOME="<path_to_the_directory_ebx_home>"
EBX_OPTIONS="-Debx.home=${EBX_HOME} -Debx.properties=${EBX_HOME}/ebx.properties"
export JAVA_OPTIONS="${EBX_OPTIONS} ${JAVA_OPTIONS}"
```

## 58.6 EBX® and third-party libraries deployment

1. Copy third-party libraries from the *EBX® CD* to the `<DOMAIN_HOME>/lib` directory. In our example, for an H2 standalone data base, we will have:

<DOMAIN\_HOME>/lib/h2-2.2.224.jar (default persistence factory) coming from ebx.software/lib/lib-h2 directory.

The complete description of the components needed by EBX® is given in chapter [Java EE deployment](#) [p 315]. Obviously, if those components are already deployed on the class-loading system, they do not have to be duplicated (ex: javax.mail-1.5.6.jar is already present in the WebLogic Server).

## 58.7 Data source and JDBC provider configuration

1. Start the 'Administration server' (for example 'AdminServer'), using:  
`<DOMAIN_HOME>/bin/startWebLogic.sh`
2. Launch the 'WebLogic Server Administration Console' by entering the following URL in the browser:  
<http://localhost:7001/console>.  
 Log in with the domain administrator username and password
3. Click on 'Services > Data sources' in the 'Domain Structure' panel, then click on 'New > Generic Data Source':
  1. Set: Type Name: EBX\_REPOSITORY, JNDI Name: EBX\_REPOSITORY, Database Type: *Your database type*  
 Click 'Next'
  2. Choose your database driver type, and click 'Next'
  3. Uncheck 'Supports Global Transactions', and click 'Next'
  4. Setup your database 'Connection Properties' and click 'Next'
  5. Click 'Test Configuration' and then 'Finish'
  6. Switch on the 'Targets' tab and select all Servers, then click 'Save'
  7. Restart the Administration server (for example 'AdminServer'), using:  
`<DOMAIN_HOME>/bin/stopWebLogic.sh`  
`<DOMAIN_HOME>/bin/startWebLogic.sh`

## 58.8 EBX® application deployment

1. Copy from the *EBX® CD* the ebx.software/webapps/ear-packaging/EBXForWebLogic.ear to the *EBX\_HOME* directory. In our example, we will have:  
`C:\EBX\home\EBXForWebLogic.ear`, or `/home/ebx/EBXForWebLogic.ear`
2. Launch the 'WebLogic Server Administration Console' by entering the following URL in the browser:  
<http://localhost:7001/console>
3. Click on 'Lock and Edit' in the 'Change Center' panel
4. Click on 'Deployments' in the 'Domain Structure' panel, and click 'Install':
  1. Install Application Assistant: Enter in 'Path' the application full path to EBXForWebLogic.ear file, located in `C:\EBX\home\`, or `/home/ebx/` directory and click 'Next'

2. Choose the installation type and scope: Click on 'Install this deployment as an application', 'Global' default scope and click 'Next'
3. Select the deployment targets: Select a node (for example 'EbxServer') from the 'Servers' list and click 'Next'
4. Optional Settings: keep as default and click 'Finish'
5. Click on 'Activate Changes', on the top left corner. The deployment status will change to 'prepared'
6. Switch to 'Control' tab, select the 'EBXForWebLogic' enterprise application, then click 'Start' > 'Servicing all requests'
7. Start the application node (for example 'EbxServer'), using:

```
<DOMAIN_HOME>/bin/startManagedWebLogic.sh EbxServer http://localhost:7001
```

## 58.9 EBX® application start

1. After WebLogic Application Server launch, run the EBX® web application by entering the following URL in the browser:  
<http://localhost:7003/ebx/>
2. At first launch, [EBX® Wizard](#) [p 369] helps to configure the default properties of the initial repository.



## CHAPTER 59

# TIBCO EBX® main configuration file

This chapter contains the following topics:

1. [Overview](#)
2. [Setting an EBX® license key](#)
3. [Setting automatic installation on first launch](#)
4. [Setting the EBX® root directory](#)
5. [Configuring the EBX® repository](#)
6. [Configuring the user and roles directory](#)
7. [Configuring EBX® localization](#)
8. [Setting temporary files directories](#)
9. [Activating the XML audit trail](#)
10. [Configuring the EBX® logs](#)
11. [Activating and configuring SMTP and emails](#)
12. [Configuring data services](#)
13. [Activating and configuring JMS](#)
14. [Configuring distributed data delivery \(D3\)](#)
15. [Configuring REST toolkit services](#)
16. [Configuring Web access from end-user browsers](#)
17. [Configuring failover](#)
18. [Tuning the EBX® repository](#)
19. [Miscellaneous](#)

## 59.1 Overview

The EBX® main configuration file, by default named `ebx.properties`, contains most of the basic parameters for running EBX®. It is a Java properties file that uses the [standard simple line-oriented format](#).

The main configuration file complements the [Java EE deployment descriptor](#) [p 319]. Administrators can also perform further configuration through the user interface, which is then stored in the EBX® repository.

See also

[Deployment details](#) [p 319]

[UI administration](#) [p 385]

## Location of the file

The access path to the main configuration file can be specified in several ways. In order of descending priority:

1. By defining the Java system property 'ebx.properties'. For example, this property can be set by adding the option `-Debx.properties=<filePath>` to the java command-line command. See [Java documentation](#).
2. By defining the servlet initialization parameter 'ebx.properties'.  
This standard Java EE setting must be specified in the `web.xml` file of the web application 'ebx'. EBX® accesses this parameter by calling the method `ServletConfig.getInitParameter("ebx.properties")` in the servlet `FrontServlet`.  
See [getInitParameter](#) in the Oracle `ServletConfig` documentation.
3. By default, if nothing is specified, the main configuration file is located at `WEB-INF/ebx.properties` of the web application 'ebx'.

### Note

In addition to specifying properties in the main configuration file, it is also possible to set the values of properties directly in the system properties. For example, using the `-D` argument of the java command-line command.

## Custom properties and variable substitution

The value of any property can include one or more variables that use the syntax `${propertyKey}`, where `propertyKey` is either a system property, or a property defined in the main configuration file.

For example, the default configuration file provided with EBX® uses the custom property `ebx.home` to set a default common directory, which is then included in other properties.

## 59.2 Setting an EBX® license key

See also [Initialization and first-launch assistant](#) [p 369]

The license key can be retrieved from the [TIBCO eDelivery](#) site.

```
#####
## EBX® License number
## (as specified by your license agreement)
#####
ebx.license=paste_here_your_license_key
```

## 59.3 Setting automatic installation on first launch

Repository can be automatically installed on first startup.

```
#####
## Installation on first launch.
## All values are ignored if the repository is already installed.
#####
```

```
## Enables repository installation on first startup (default is false).
## If true, property ebx.license should also be set to a valid license.
ebx.install.enabled=true

## Following properties configure the repository. Values are optional and defaults are automatically generated.
ebx.install.repository.id=00275930BB88
ebx.install.repository.label=A Test

## Following properties specify the EBX administrator. These are ignored if a custom directory is defined.
ebx.install.admin.login=admin
ebx.install.admin.firstName=admin
ebx.install.admin.lastName=admin
ebx.install.admin.email=adamin@example.com

## Following property specifies the none encrypted password used for the EBX administrator.
## It is ignored if a custom directory is defined. It cannot be set if property
  ebx.install.admin.password.encrypted is set.
#ebx.install.admin.password=admin

## Following property specifies the encrypted password used for the EBX administrator.
## It is ignored if a custom directory is defined. It cannot be set if property ebx.install.admin.password is
  set.
## Password can be encrypted by using command:
## java -cp ebx.jar com.orchestranetworks.service.directory.EncryptPassword password_to_encrypt
ebx.install.admin.password.encrypted=8c6976e5b5410415bde908bd4dee15dfb167a9c873fc4bb8a81f6f2ab448a918
```

## 59.4 Setting the EBX® root directory

The EBX® root directory contains archives, the XML audit trail and, when the repository is persisted on H2 standalone mode, the H2 database files.

```
#####
## Path for EBX® XML repository
#####
ebx.repository.directory=${ebx.home}/ebxRepository
```

## 59.5 Configuring the EBX® repository

Before configuring the persistence properties of the EBX® repository, carefully read the section [Technical architecture](#) [p 374] in the chapter 'Repository administration'.

The required library (driver) for each supported database is described in the chapter [Database drivers](#) [p 316].

### See also

[Repository administration](#) [p 374]

[Rules for the database access and user privileges](#) [p 375]

[Supported databases](#) [p 311]

[Data source of the EBX® repository](#) [p 320]

[Database drivers](#) [p 316]

```
#####
## The maximum time to set up the database connection,
## in milliseconds.
#####
ebx.persistence.timeout=10000
#####
## The prefix to add to all table names of persistence system.
## This may be useful for supporting multiple repositories in the relational database.
## Default value is 'EBX_'.
#####
ebx.persistence.table.prefix=

#####
## Case EBX® persistence system is H2 'standalone'.
#####
ebx.persistence.factory=h2.standalone
ebx.persistence.user=sa
```

```

ebx.persistence.password=

#####
## Case EBX® persistence system is H2 'server mode',
#####
#ebx.persistence.factory=h2.server

## Specific properties to be set only only if you want to ignore the standard
## deployment process of 'ebx' web application in the target operational environment
## (see the deployment descriptor 'web.xml' of 'ebx' web application).
#ebx.persistence.url=jdbc:h2:tcp://127.0.0.1/ebxdb
#ebx.persistence.user=xxxxxxx
#ebx.persistence.password=yyyyyyy

#####
## Case EBX® persistence system is Oracle database.
#####
#ebx.persistence.factory=oracle

## Specific properties to be set only only if you want to ignore the standard
## deployment process of 'ebx' web application in the target operational environment
## (see the deployment descriptor 'web.xml' of 'ebx' web application).
#ebx.persistence.url=jdbc:oracle:thin:@127.0.0.1:1521:ebxDatabase
#ebx.persistence.driver=oracle.jdbc.OracleDriver
#ebx.persistence.user=xxxxxxx
#ebx.persistence.password=yyyyyyy

## Activate to use VARCHAR2 instead of NVARCHAR2 on Oracle; never modify on an existing repository.
#ebx.persistence.oracle.useVARCHAR2=false

#####
## Case EBX® persistence system is Microsoft SQL Server.
#####
#ebx.persistence.factory=sqlserver

## Specific properties to be set only only if you want to ignore the standard
## deployment process of 'ebx' web application in the target operational environment
## (see the deployment descriptor 'web.xml' of 'ebx' web application).
#ebx.persistence.url= \
#jdbc:sqlserver://127.0.0.1:1036;datasenname=ebxDatabase
#ebx.persistence.driver=com.microsoft.sqlserver.jdbc.SQLServerDriver
#ebx.persistence.user=xxxxxxx
#ebx.persistence.password=yyyyyyy

#####
## Case EBX® persistence system is Microsoft Azure SQL database.
#####
#ebx.persistence.factory=azure.sql

## Specific properties to be set only only if you want to ignore the standard
## deployment process of 'ebx' web application in the target operational environment
## (see the deployment descriptor 'web.xml' of 'ebx' web application).
#ebx.persistence.url= \
#jdbc:sqlserver://myhost.database.windows.net:1433;database=ebxDatabase;encrypt=true;\
#trustServerCertificate=false;hostNameInCertificate=*.database.windows.net;
#ebx.persistence.driver=com.microsoft.sqlserver.jdbc.SQLServerDriver
#ebx.persistence.user=xxxxxxx
#ebx.persistence.password=yyyyyyy

#####
## Case EBX® persistence system is PostgreSQL.
#####
#ebx.persistence.factory=postgresql

## Specific properties to be set only only if you want to ignore the standard
## deployment process of 'ebx' web application in the target operational environment
## (see the deployment descriptor 'web.xml' of 'ebx' web application).
#ebx.persistence.url=jdbc:postgresql://127.0.0.1:5432/ebxDatabase
#ebx.persistence.driver=org.postgresql.Driver
#ebx.persistence.user=xxxxxxx
#ebx.persistence.password=yyyyyyy

```

## 59.6 Configuring the user and roles directory

This parameter specifies the Java directory factory class name. It must only be defined if not using the default EBX® directory.



## See also

[Users and roles directory](#) [p 401]

*DirectoryFactory*<sup>API</sup>

```
#####
## Specifies the Java directory factory class name.
## Value must be the fully qualified name of the Java class.
## The class must extend com.orchestranetworks.service.directory.DirectoryFactory.
#####
#ebx.directory.factory=xxx.yyy.DirectoryFactoryImpl
```

It is also possible to disable the built-in role "ADMINISTRATOR".

```
#####
## Specifies whether the built-in role ADMINISTRATOR is disabled.
## Default value is false.
#####
#ebx.directory.disableBuiltInAdministrator=true
```

## 59.7 Configuring EBX® localization

This parameter is used to configure the locales used at runtime. This list must contain all the locales that are exposed to the end-user. EBX® will not be able to display labels and messages in a language that is not declared in this list.

The default locale must be the first one in the list.

```
#####
## Available locales, separated by a comma.
## The first element in the list is considered as the default locale.
## If not set, available locales are 'en-US, fr-FR'.
##
#####
#ebx.locales.available=en-US, fr-FR
```

See also [Extending TIBCO EBX® internationalization](#) [p 237]

## 59.8 Setting temporary files directories

Temporary files are stored as follows:

```
#####
## Directories for temporary resources.
#####
# The property ebx.temp.directory allows to specify a directory for temporary files.
# Default value is java.io.tmpdir
#
#ebx.temp.directory = \\${java.io.tmpdir}
#ebx.temp.directory = /tmp/java

# The property ebx.temp.cache.directory allows to specify the directory containing temporary files for cache.
# Default value is ${ebx.temp.directory}/ebx.platform.
#ebx.temp.cache.directory = ${ebx.temp.directory}/ebx.platform

# The property ebx.temp.import.directory allows to specify the directory containing temporary files for import.
# Default value is ${ebx.temp.directory}/ebx.platform.
#ebx.temp.import.directory = ${ebx.temp.directory}/ebx.platform
```

## 59.9 Activating the XML audit trail

By default, the XML audit trail is activated. It can be deactivated using the following variable:

```
#####
# The XML history has been replaced by an SQL history.
# This old XML history can be deactivated using the following variable.
# Default is true.
```

```
#####
ebx.history.xmlaudittrail.activated = true
```

See also [Audit trail](#) [p 421]

## 59.10 Configuring the EBX® logs

The most important logging categories are:

|                                            |                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ebx.log4j.category.log.kernel</b>       | Logs for EBX® main features, processes, exceptions and compilation results of modules and data models.                                                                                                                                                                              |
| <b>ebx.log4j.category.log.workflow</b>     | Logs for main features, warnings and exceptions about workflow.                                                                                                                                                                                                                     |
| <b>ebx.log4j.category.log.persistence</b>  | Logs related to communication with the underlying database.                                                                                                                                                                                                                         |
| <b>ebx.log4j.category.log.setup</b>        | Logs for the compilation results of all EBX® objects, except for modules and data models.                                                                                                                                                                                           |
| <b>ebx.log4j.category.log.validation</b>   | Logs for datasets validation results.                                                                                                                                                                                                                                               |
| <b>ebx.log4j.category.log.mail</b>         | <p>Logs for the activity related to the emails sent by the server (see <a href="#">Activating and configuring SMTP and emails</a> [p 357]).</p> <p><b>Note:</b> This category must not use the <a href="#">Custom SMTP appender</a> [p 355] in order to prevent infinite loops.</p> |
| <b>ebx.log4j.category.log.d3</b>           | Logs for D3 events on EBX®.                                                                                                                                                                                                                                                         |
| <b>ebx.log4j.category.log.dataservices</b> | Logs for data service events in EBX®.                                                                                                                                                                                                                                               |
| <b>ebx.log4j.category.log.monitoring</b>   | Raw logs for <a href="#">monitoring</a> [p 297].                                                                                                                                                                                                                                    |
| <b>ebx.log4j.category.log.request</b>      | The optimization strategy for every Request <sup>API</sup> issued on a semantic table in the EBX® repository.                                                                                                                                                                       |
| <b>ebx.log4j.category.log.restServices</b> | Logs for REST services events in EBX®, including those from the <a href="#">REST Toolkit</a> [p 721].                                                                                                                                                                               |

Some of these categories can also be written to through custom code using the `LoggingCategoryAPI` interface.

```
#####
```

```

## Log4J properties:
##
## We have some specific syntax extensions:
## - Appender ebxFile:<aFileName>
## Defines a file appender with default settings (threshold=DEBUG)
##
## - property log.defaultConversionPattern is set by Java
#####
#ebx.log4j.debug=true
#ebx.log4j.disable=
ebx.log4j.rootCategory= INFO
ebx.log4j.category.log.kernel= INFO, Console, ebxFile:kernel, kernelMail
ebx.log4j.category.log.workflow= INFO, ebxFile:workflow
ebx.log4j.category.log.persistence= INFO, ebxFile:persistence
ebx.log4j.category.log.setup= INFO, Console, ebxFile:kernel
ebx.log4j.category.log.mail= INFO, Console, ebxFile:mail
ebx.log4j.category.log.frontEnd= INFO, Console, ebxFile:kernel
ebx.log4j.category.log.frontEnd.incomingRequest= INFO
ebx.log4j.category.log.frontEnd.requestHistory= INFO
ebx.log4j.category.log.frontEnd.UIComponentInput= INFO
ebx.log4j.category.log.fsm= INFO, Console, ebxFile:fsm
ebx.log4j.category.log.fsm.dispatch= INFO
ebx.log4j.category.log.fsm.pageHistory= INFO
ebx.log4j.category.log.wbp= FATAL, Console
#-----
ebx.log4j.appender.Console.Threshold = INFO
ebx.log4j.appender.Console=com.onwbp.org.apache.log4j.ConsoleAppender
ebx.log4j.appender.Console.layout=com.onwbp.org.apache.log4j.PatternLayout
ebx.log4j.appender.Console.layout.ConversionPattern=${log.defaultConversionPattern}
#-----
ebx.log4j.appender.kernelMail.Threshold = ERROR
ebx.log4j.appender.kernelMail = com.onwbp.org.apache.log4j.net.SMTPAppender
ebx.log4j.appender.kernelMail.To = admin@domain.com
ebx.log4j.appender.kernelMail.From = admin${ebx.site.name}
ebx.log4j.appender.kernelMail.Subject = EBX® Error on Site ${ebx.site.name} (VM ${ebx.vm.id})
ebx.log4j.appender.kernelMail.layout.ConversionPattern=**Site ${ebx.site.name} (VM${ebx.vm.id})*%n
${log.defaultConversionPattern}
ebx.log4j.appender.kernelMail.layout = com.onwbp.org.apache.log4j.PatternLayout
#-----
ebx.log4j.category.log.monitoring= INFO, ebxFile:monitoring
ebx.log4j.category.log.dataServices= INFO, ebxFile:dataServices
ebx.log4j.category.log.d3= INFO, ebxFile:d3
ebx.log4j.category.log.request= INFO, ebxFile:request
ebx.log4j.category.log.restServices= INFO, ebxFile:dataServices

```

## Custom 'ebxFile' appender

The token ebxFile: can be used as a shortcut to define a daily rolling file appender with default settings. It must be followed by a file name. It then activates an appender that writes to a file located in the directory ebx.logs.directory.

The property ebx.log4j.appender.ebxFile.backup.Threshold allows defining the maximum number of backup files for daily rollover.

```

#####
## Directory of log files 'ebxFile:'
## This property is used by special appender prefixed
## by 'ebxFile:' (see log section below)
#####
ebx.logs.directory=${ebx.home}/ebxLog

#####
# Daily rollover threshold of log files 'ebxFile:'
# Specifies the maximum number of backup files for daily rollover of 'ebxFile:' appenders.
# When set to a negative value, backup log files are never purged.
# Default value is -1.
#####
ebx.log4j.appender.ebxFile.backup.Threshold=-1

```

## Custom SMTP appender

The appender com.onwbp.org.apache.log4j.net.SMTPAppender provides an asynchronous email sender.

See also [Activating and configuring SMTP and emails](#) [p 357]

## Module's appenders configuration

The configurations of module logging categories are declared according to a hierarchical pattern. Furthermore, every enabled logging request for a given category will be forwarded to all the appenders defined for that category, as well as to the appenders higher in the hierarchy. This fact means that appenders are inherited additively.

The root module logging category can be customized by setting the property `ebx.log4j.category.log.wbp`. For a given module, if the root module logging category configuration is the only one applicable, then a `DailyRollingFileAppender` is automatically added to the module's appenders list. The log file name will be derived from the module's name and from an optional specific sub-category.

```
#####
## Root module logging category configuration
#####
ebx.log4j.category.log.wbp = FATAL, Console
```

Every module logging category can be customized by setting the property `ebx.log4j.category.log.module.xxxxxx`, where `xxxxxx` corresponds to the module's name. The explicit configuration of the module logging category will disable the automatic addition of the `DailyRollingFileAppender` previously described. However, any other defined appenders will be inherited additively.

Since the inheritance mechanism may not be suitable for every case, the additivity can be broken by setting to `false` the property `ebx.log4j.additivity.log.wbp.xxxxxx`, where `xxxxxx` corresponds to the module's name.

```
#####
## Module logging category configuration
##
## The module's log messages will only be written to mycompany-module log file
## since additivity has been broken
#####
ebx.log4j.category.log.module.mycompany-module = INFO, ebxFile:mycompany-module

ebx.log4j.additivity.log.wbp.mycompany-module = false
```

## Module's log threshold

The inheritance mechanism described in [Module's appenders configuration](#) [p 356] is applied to the module log threshold as well. Actually, the inherited level for a given logger, is equal to the first non-null level in the hierarchy, starting from this logger and proceeding up to the root module logging category.

## Custom module log threshold

There is an exception to the inheritance, for custom module, since the log level threshold of their logging category, by default, is set to `INFO`. This threshold can be customized by setting the property `ebx.log4j.category.log.module.xxxxxx`, where `xxxxxx` corresponds to the custom module's name.

Example: `ebx.log4j.category.log.module.mycompany-module=DEBUG`.

See also `ModuleContextOnRepositoryStartup.getLoggingCategory`<sup>API</sup>

## Add-on module log threshold

Like custom module, by default, the log level threshold of any add-on module is set to INFO.

The log level threshold can be customized by setting the property `ebx.log4j.category.log.addon.xxxxxx` where `xxxxxx` corresponds to the add-on module's name.

Example: `ebx.log4j.category.log.addon.daqa=DEBUG`

## Modules log threshold overrides summary

Thus and considering EBX® logging features, the log threshold defined at the root level in the logger hierarchy may be overridden by (by order from least to most weighted):

- the module logging category explicit configuration,
- the module's `log.threshold.xxxx` or `log.threshold.category.xxxx` log configuration properties, where `xxxx` corresponds to the sub category,
- the following log configuration properties: `ebx.log4j.category.log.module.xxxxxx`, where `xxxxxx` corresponds to the custom module's name, or `ebx.log4j.category.log.addon.xxxxxx`, where `xxxxxx` corresponds to the add-on module's name.

## 59.11 Activating and configuring SMTP and emails

The internal mail manager sends emails asynchronously. It is used by the workflow engine and the custom SMTP appender `com.onwbp.org.apache.log4j.net.SMTPAppender`.

See also [Mail sessions](#) [p 321]

```
#####
## SMTP and emails
#####

## Activate emails (true or false, default is false).
## If activated, the deployer must ensure that the entry 'mail/EBX_MAIL_SESSION' is bound
## in the operational environment of the application server (except if a specific email
## configuration is used by setting the property ebx.mail.smtp.host below).
#ebx.mail.activate=false

## Polling interval is in seconds (default is 10).
#ebx.mail.polling.interval=10

## Specific properties to be set only only if you want to ignore the standard
## deployment process of 'ebx' web application in the target operational environment
## (see the deployment descriptor 'web.xml' of 'ebx' web application).
#ebx.mail.smtp.host = smtp.domain.com
## SMTP port default is 25.
#ebx.mail.smtp.port= 25
#ebx.mail.smtp.login=
#ebx.mail.smtp.password=
## Activate SSL (true or false, default is false).
#ebx.mail.smtp.ssl.activate=true
```

## 59.12 Configuring data services

```
#####
## Data services
#####

# Specifies the default value of the data services parameter
# 'disableRedirectionToLastBroadcast'.
# Default is false.
#ebx.dataservices.disableRedirectionToLastBroadcast.default=false

# Specifies the default value for deletion at the end of close and
```

```
# merge operations.
# If the parameter is set in the request operation, it overrides
# this default setting.
# If unspecified, default is false.
#ebx.dataservices.dataDeletionOnCloseOrMerge.default=false
#ebx.dataservices.historyDeletionOnCloseOrMerge.default=false

# Upon WSDL generation, specifies if the target namespace value
# corresponds to the content before 5.5.0 'ebx-services'
# or 'urn:ebx:ebx-services' in conformity with the URI syntax.
# If the parameter is set to true, there is no check of the target
# namespace as URI at the WSDL generation.
# If unspecified, default is false.
#ebx.dataservices.wsdlTargetNamespace.disabledCheck=false

#####
## REST configuration
#####

# If activated, the HTTP request header 'Accept' is used to specify
# the accepted content type. If none is supported, an error is
# returned to the client with the HTTP code 406 'Not acceptable'.
# If deactivated, the header is ignored therefore the best content
# type is used.
# Default is false.
#ebx.dataservices.rest.request.checkAccept=false

# If activated, it tries authentication 'Basic Authentication Scheme'
# method and set 'Basic' value in 'WWW-Authenticate' header of HTTP
# response.
# Default is false.
#ebx.dataservices.rest.auth.tryBasicAuthentication=false

# Authorization token timeout is seconds.
# Default value is 1800 seconds (30 minutes)
# This value is ignored if 'Token Authentication Scheme' is not activated.
#ebx.dataservices.rest.auth.token.timeout=1800
```

## 59.13 Activating and configuring JMS

See also [JMS for data services](#) [p 322]

```
#####
## JMS configuration for Data Services
#####

## Activates JMS (true or false, default is false).
## If activated, the deployer must ensure that the entry 'jms/EBX_JMSConnectionFactory'
## are bound in the operational environment of the application server.
## The entry 'jms/EBX_QueueIn' should also be bound to enable handling Data Services
## request using JMS.
#ebx.jms.activate=false

## Activates JMS queue for failures (true or false, default is false).
## If activated, the deployer must ensure that the entry 'jms/EBX_QueueFailure' is bound
## in the operational environment of the application server.
#ebx.jms.activate.queueFailure=false

## Number of concurrent listener(s)
## Default is 3.
## Property is used if ebx.jms.activate is set to true.
#ebx.jms.listeners.count=3
```

## 59.14 Configuring distributed data delivery (D3)

See [Configuring D3 nodes](#) [p 445] for the main configuration file properties pertaining to D3.

See also

[JMS for distributed data delivery \(D3\)](#) [p 435]

[Introduction to D3](#) [p 426]

# 59.15 Configuring REST toolkit services

```
#####
## REST configuration
#####

# Defines the maximum number of bytes that will be extracted
# from the REST request body to build some DEBUG log messages.
# Default value is 8192 bytes.
# This value is ignored if DEBUG level is not activated on the restServices logger.
#ebx.restservices.log.body.content.extract.size=8192
```

# 59.16 Configuring Web access from end-user browsers

## HTTP Authorization header policy

EBX® natively offers three policies to send and receive credentials using HTTP headers:

|                 |                                                                                                                                                                 |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>standard</b> | It corresponds to the authentication scheme, using the HTTP Authorization header, described in the <a href="#">RFC 2617</a> .                                   |
| <b>ebx</b>      | To prevent HTTP Authorization header override issues, this policy acts the same as the standard but the credentials are stored in an EBX® specific HTTP header. |
| <b>both</b>     | It is the combination of the two previously described policies.                                                                                                 |

```
#####
## EBX® authorization header policy for HTTP requests
##
## Possible values are: standard, ebx, both.
## standard:
##   the standard HTTP Authorization header holds the credentials
## ebx:
##   an EBX® specific HTTP header holds the credentials
## both:
##   both (standard and specific) HTTP headers hold the credentials
##
## Default value is: both.
#####
#ebx.http.authorization.header.policy=both
```

## URLs computing

By default, EBX® runs in "standalone" mode, where external resources (images, JavaScript, etc.) are provided by the application server.

Also by default, URL-related parameters in the main configuration file do not have to be set. In this case, the server name and the port are obtained from the initial request sent to EBX®.

See also [URL policy \(deprecated\)](#) [p 388]

```
#####
## EBX® FrontServlet: default properties for computing servlet address
##
## {useLocalUrl}:
## If set to true, servlet address is a "local absolute" URL.
## (that is, a relative URL consisting of an absolute path: "/path")
```

```

## See RFC 2396, http://www.ietf.org/rfc/rfc2396.txt).
## This property is defined once for HTTP and HTTPS.
## Default value is false.
##
## {host}:
## If neither defined nor adapted, retrieves initial request host
## {port}:
## If neither defined nor adapted, retrieves initial request host
## {path}:
## Mandatory, may be empty
## {ui.path}:
## If not defined, defaults to ebx-ui/
## {http.useHttpsSettings}:
## If true, force the use of SSL security even if the incoming
## requests do not. Default value is false.
##
## Resulting address will be:
## EBX®: protocol://{host}:{port}/{path}
## UI: protocol://{host}:{port}/{ui.path}
##
## Each property for HTTP (except {port}) may be inherited from HTTPS property,
## and reciprocally.
#####

ebx.servlet.useLocalUrl=true

#ebx.servlet.http.host=
#ebx.servlet.http.port=
ebx.servlet.http.path=ebx/
#ebx.servlet.http.ui.path=ebx-ui/
#ebx.servlet.http.useHttpsSettings=false

#ebx.servlet.https.host=
#ebx.servlet.https.port=
ebx.servlet.https.path=ebx/
#ebx.servlet.https.ui.path=ebx-ui/

#####
## External resources: default properties for computing external resources address
##
## The same rules apply as EBX® FrontServlet properties (see comments).
##
## Each property may be inherited from EBX® FrontServlet.
#####

ebx.externalResources.useLocalUrl=true

#ebx.externalResources.http.host=
#ebx.externalResources.http.port=
#ebx.externalResources.http.path=
#ebx.externalResources.http.useHttpsSettings=false

#ebx.externalResources.https.host=
#ebx.externalResources.https.port=
#ebx.externalResources.https.path=

```

## Proxy mode

Proxy mode allows using a front-end HTTP server to provide static resources (images, CSS, JavaScript, etc.). This architecture reduces the load on the application server for static HTTP requests. This configuration also allows using SSL security on the front-end server.

The web server sends requests to the application server according to a path in the URL. The `servletAlias` and `uiServletAlias` paths are specified in the main configuration file.

The web server provides all external resources. These resources are stored in a dedicated directory, accessible using the `resourcesAlias` path.

EBX® must also be able to access external resources from the file system. To do so, the property `ebx.webapps.directory.externalResources` must be specified.

To force the use of SSL security even if the incoming requests do not, `ebx.servlet.http.useHttpsSettings` and / or `ebx.externalResources.http.useHttpsSettings` properties must be set to true. Their default values are false.



The main configuration file may be configured as follows:

```
#####
## Path for external resources if they are not
## delivered within web applications
## This field is mandatory if in proxy mode.
#####
ebx.webapps.directory.externalResources=D:/http/resourcesFolder

#####

ebx.servlet.useLocalUrl=true

#ebx.servlet.http.host=
#ebx.servlet.http.port=
ebx.servlet.http.path=servletAlias
ebx.servlet.http.ui.path=uiServletAlias
#ebx.servlet.http.useHttpsSettings=false

#ebx.servlet.https.host=
#ebx.servlet.https.port=
ebx.servlet.https.path=servletAlias
ebx.servlet.https.ui.path=uiServletAlias

#####

ebx.externalResources.useLocalUrl=true

#ebx.externalResources.http.host=
#ebx.externalResources.http.port=
ebx.externalResources.http.path=resourcesAlias
#ebx.externalResources.http.useHttpsSettings=false

#ebx.externalResources.https.host=
#ebx.externalResources.https.port=
ebx.externalResources.https.path=resourcesAlias
```

### Attention

When proxy mode is used, the URL to the ebx-dataservices module must be configured through the lineage administration panel. Note that the provided URL must end its path with /ebx-dataservices.

See also [Path recommendations](#)

## Reverse-proxy mode

If URLs generated by EBX®, for requests and external resources, must contain a different protocol than the one from the incoming request, a specific server name, a specific port number or a specific path prefix, properties may be configured as follows:

```
#####
#ebx.servlet.useLocalUrl=false

ebx.servlet.http.host=reverseDomain
#ebx.servlet.http.port=
ebx.servlet.http.path=ebx/
#ebx.servlet.http.ui.path=ebx-ui/
#ebx.servlet.http.useHttpsSettings=false

ebx.servlet.https.host=reverseDomain
#ebx.servlet.https.port=
ebx.servlet.https.path=ebx/
#ebx.servlet.https.ui.path=ebx-ui/

#####
## Web parameters (for external resources)
## if nothing is set, values are taken from servlet.
#####
#ebx.externalResources.useLocalUrl=false

#ebx.externalResources.http.host=
#ebx.externalResources.http.port=
#ebx.externalResources.http.path=
ebx.externalResources.http.useHttpsSettings=true
```

```
ebx.externalResources.https.host=reverseDomain
#ebx.externalResources.https.port=
ebx.externalResources.https.path=
```

### Attention

When reverse-proxy mode is used, the URL to the ebx-dataservices module must be configured through the lineage administration panel. Note that the provided URL must end its path with /ebx-dataservices.

See also [Path recommendations](#)

## Path recommendations

It is recommended to deploy webapps of EBX® modules on the same path level. For example:

- /module1
- /module2

or:

- /path/module1
- /path/module2

It avoids difficulties when configuring the advanced features of [Proxy mode](#) and [Reverse-proxy mode](#).

Consequently, it is not recommended to use the root path "/" as the servlet's HTTP or HTTPS path in the following properties:

- ebx.servlet.http.path
- ebx.servlet.http.ui.path
- ebx.servlet.https.path
- ebx.servlet.https.ui.path

## 59.17 Configuring failover

These parameters are used to configure the failover mode and activation key, as well as heartbeat logging in DEBUG mode.

See also [Failover with hot-standby](#) [p 375]

```
#####
## Mode used to qualify the way in which a server accesses the repository.
## Possible values are: unique, failovermain, failoverstandby.
## Default value is: unique.
#####
#ebx.repository.ownership.mode=unique

## Activation key used in case of failover. The backup server must include this
## key in the HTTP request used to transfer exclusive ownership of the repository.
## The activation key must be an alphanumeric ASCII string longer than 8 characters.
#ebx.repository.ownership.activationkey=

## Specifies whether to hide heartbeat logging in DEBUG mode.
## Default value is true.
#ebx.repository.ownership.hideHeartBeatLogForDebug=true
```

## 59.18 Tuning the EBX® repository

Some options can be set so as to optimize memory usage.

The properties are configured as follows:

```
#####
## Technical parameters for memory and performance tuning
#####
# Import commit threshold allows to specify the commit threshold
# exclusively for the archive import launched directly from Manager.
#
# For more details about the commit threshold,
# see the Javadoc ProcedureContext.setCommitThreshold().
# Default value is 0.
#
ebx.manager.import.commit.threshold=100

# A validation messages threshold allows specifying the maximum number of
# messages to consider per constraint when performing a validation.
# This threshold is considered for each constraint defined in a data model
# and for each severity in each dataset validation report.
# When the threshold is reached by a constraint and a severity:
# - the validation of the constraint is stopped
# - an error message indicating that the threshold has been reached
# is added to the validation report.
# When set to 0 or a negative value, the threshold is not considered.
# Default value is 0.
#
ebx.validation.constraints.messages.threshold = 100

# Specifies whether the validation report should be kept in memory,
# regardless of the loading strategy of the dataspace.
# Default value is true. However, it is recommended to deactivate it
# when the repository contains a large number of open dataspace and
# datasets.
#ebx.validation.report.keepInMemory=false
```

See also [Validation report page](#) [p 368]

## 59.19 Miscellaneous

### *Activating data workflows*

This parameter specifies whether data workflows are activated. This parameter is not taken into account on the fly. The server must be restarted whenever the value changes.

```
#####
## Workflow activation.
## Default is false.
#####
ebx.workflow.activation = true
```

### *Disabling user task legacy mode*

This parameter specifies whether the creation service of a user task in legacy mode should be offered in the workflow modeling. The default value is true.

See `UserTask.UserTaskMode.LEGACY_MODEAPT` for more information.

```
## Disables legacy work item mode(default is true)
## Specify if the creation service of user task in legacy mode must be offered
## in workflow modeling.
#ebx.manager.workflow.legacy.userTaskMode=false
```

## ***Log procedure starts***

This parameter specifies whether starts of the procedure execution are logged.

```
#####
## Specifies whether transaction starts are logged. Default is false.
#####
ebx.logs.logTransactionStart = true
```

## ***Log validation starts***

This parameter specifies whether starts of datasets validation are logged.

```
#####
## Specifies whether validation starts are logged. Default is false.
#####
ebx.logs.logValidationStart = true
```

## ***Deployment site identification***

This parameter allows specifying the email address to which technical log emails are sent.

```
#####
## Unique Site Name
## --> used by monitoring emails and by the repository
#####
ebx.site.name= name@domain.com
```

## ***Dynamically reloading the main configuration***

Some parameters can be dynamically reloaded, without restarting EBX®. The parameter `thisfile.checks.intervalInSeconds` indicates how frequently the main configuration file is checked.

```
#####
### Checks if this file has been updated
### If value <= 0, no more checks will be done
#####
thisfile.checks.intervalInSeconds=1
```

In development mode, this parameter can be set to as low as one second. On production systems, where changes are expected to be less frequent, the value can be greater, or set to '0' to disable hot reloading entirely.

This property is not always supported when the module is deployed as a WAR, as it would then depend on the application server.

## ***Declaring modules as undeclared***

On application server startup, the initialization of deployed web applications / EBX® modules and the initialization of the EBX® repository are performed asynchronously. In order to properly initialize the EBX® repository, it is necessary to compile all the data models used by at least a dataset, hence EBX® will wait endlessly for referenced modules to be registered.

If a module is referenced by a data model but is not deployed (or no longer deployed), it is necessary to declare this module as undeployed to unlock the wait and continue the startup process.

#### Note

The `kernel` logging category indicates which modules are awaited.

#### Note

A module declared as undeployed cannot be registered into EBX® until it is removed from the property `ebx.module.undeployedModules`.

#### Note

Any data model based on an unregistered module will have an "undeployed module" compilation error.

See also

[Module registration](#) [p 462]

[Dynamically reloading the main configuration](#) [p 364]

```
#####
## Comma-separated list of EBX® modules declared
## as undeployed.
## If a module is expected by the EBX® repository but is
## not deployed, it must be declared in this property.
## Caution:
## if the "thisfile.checks.intervalInSeconds" property is deactivated,
## a restart is mandatory, otherwise it will be hot-reloaded.
#####
ebx.module.undeployedModules=
```

## Module public path prefix

EBX® modules' public paths are declared in the 'module.xml' file of each module. A context prefix can be declared for all modules, without having to modify the 'module.xml' content, by specifying the property that follows.

This prefix will apply to any EBX® module, including core, add-on and specific modules.

When proxy and / or reverse-proxy mode are used, the `ebx.servlet.http[s].path` and `ebx.servlet.http[s].ui.path` properties must take into account this module public path prefix setting. Conversely, the `ebx.externalResources.http[s].path` property must end its path just before a potential prefix.

```
#####
ebx.servlet.useLocalUrl=true

#ebx.servlet.http.host=
#ebx.servlet.http.port=
ebx.servlet.http.path=reverse-proxy/prefix/ebx/
ebx.servlet.http.ui.path=reverse-proxy/prefix/ebx-ui/
#ebx.servlet.http.useHttpsSettings=false

#ebx.servlet.https.host=
#ebx.servlet.https.port=
ebx.servlet.https.path=reverse-proxy/prefix/ebx/
ebx.servlet.https.ui.path=reverse-proxy/prefix/ebx-ui/

#####
## Web parameters (for external resources)
## if nothing is set, values are taken from servlet.
#####
ebx.externalResources.useLocalUrl=true

#ebx.externalResources.http.host=
#ebx.externalResources.http.port=
```

```

ebx.externalResources.http.path=reverse-proxy/
#ebx.externalResources.http.useHttpsSettings=false

#ebx.externalResources.https.host=
#ebx.externalResources.https.port=
ebx.externalResources.https.path=reverse-proxy/

#####
## EBX® Module context path prefix
##
## If defined, applies to all EBX® modules public paths declared in
## any module.xml file (core, add-on and specific).
#####
ebx.module.publicPath.prefix=prefix/

```

See [URLs computing](#) [p 359] for more information.

## **EBX® run mode**

This property defines how EBX® runs. Three run modes are available: *development*, *integration* and *production*.

When running in *development* mode, the [development tools](#) [p 473] are activated in EBX®, some features thus become fully accessible and more technical information is displayed.

### **Note**

The administrator can always access this information regardless of the mode used.

The additional features accessible when running in *development* mode include the following (non-exhaustive list):

|                                     |                                                                                                                                                                                                            |
|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Documentation pane</b>           | In the case of a computed value, the Java class name is displayed. A button is displayed giving access to the path to a node.                                                                              |
| <b>Compilation information</b>      | Module and schema compilation information is displayed in the dataset validation report.                                                                                                                   |
| <b>Java bindings</b>                | The generation of Java bindings is available if the schema of the dataset mentions at least one binding.                                                                                                   |
| <b>Web component link generator</b> | The Web component link generator is available on datasets and dataspace.                                                                                                                                   |
| <b>Data model assistant</b>         | Data model configuration and additional options, such as Services, Business Objects and Rules, Java Bindings, Toolbars and some advanced properties.                                                       |
| <b>Workflow modeling</b>            | Declare specific script tasks.                                                                                                                                                                             |
| <b>Log</b>                          | The logs include additional technical information intended for the developer. For example, a warning is written to logs if a drop-down list is defined on a node which is not an enumeration in a UI Bean. |
| <b>Product documentation</b>        | The product documentation is always the most complete one (i.e "advanced"), including administration and development chapters.                                                                             |

```
#####
## Server Mode
## Value must be one of: development, integration, production
## Default is production.
#####
backend.mode=integration
```

#### Note

There is no difference between the *integration* and *production* modes.

## Resource filtering

This property allows the filtering of certain files and directories in the resource directory contents (resource type node, with an associated facet that indicates the directory that contains usable resources).

```
#####
## list (separated by comma) of regexps excluding resource
## the regexp must be of type "m:[pattern]:[options]".
## the list can be void
#####
```

```
ebx.resource.exclude=m:CVS/*:
```

## Validation report page

The validation report page can display a finite number of items for each severity. This number can be tuned with this property.

```
# Defines the maximum item displayed for each severity in the validation report page.  
# Default value is 100.  
#ebx.validation.report.maxItemDisplayed=100
```

See also [Tuning the EBX® repository](#) [p 363]

## Validation report logs

This property allows to specify the number of validation messages to display in the logs when validating a dataset or a table.

```
# Defines the maximum number of messages displayed in the logs.  
# Default value is 100.  
# When set to 0 or a negative value, the limit is not considered.  
#ebx.validation.report.maxItemDisplayedInLogs=500
```

See also [Tuning the EBX® repository](#) [p 363]



## CHAPTER 60

# Initialization and first-launch assistant

Deliverables can be found on [TIBCO eDelivery](#) (an account is mandatory in order to access eDelivery, please contact [the support team](#) to request one).

The TIBCO EBX® Configuration Assistant helps with the initial configuration of the EBX® repository. If EBX® does not have a repository installed upon startup and if the [automatic installation](#) [p 350] is not enabled, the configuration assistant is launched automatically.

Before starting the configuration of the repository, make sure that EBX® is correctly deployed on the application server. See [Java EE deployment](#) [p 315].

### Note

The EBX® main configuration file must also be properly configured. See [TIBCO EBX® main configuration file](#) [p 349].

This chapter contains the following topics:

1. [License key](#)
2. [Configuration steps](#)

## 60.1 License key

When launching EBX®, the license key page displays automatically if no valid license key is available, that is, if there is no license key entered in the EBX® main configuration file, or if the current license key has expired.

The license key can be retrieved from the [TIBCO eDelivery](#) site (an account is mandatory in order to access eDelivery, please contact <https://support.tibco.com> to request one).

The license key is available in the `TIB_ebx_5.9.26_license_key.pdf` file.

## 60.2 Configuration steps

The EBX® configuration assistant guides you through the following steps:

1. Validating the license agreement.
2. Configuring the repository.

3. Defining users in the default user and roles directory (if a custom directory is not defined).
4. Validating the information entered.
5. Installing the EBX® repository.

## ***Validating the license agreement***

In order to proceed with the configuration, you must read and accept the product license agreement.

## ***Configuring the repository***

This page displays some of the properties defined in the EBX® main configuration file. You also define several basic properties of the repository in this step.

---

|                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Id of the repository</b><br>( <code>repositoryId</code> ) | Must uniquely identify the repository (in the scope of the enterprise). The identifier is 48 bits (6 bytes) long and is usually represented as 12 hexadecimal digits. This information is used for generating the Universally Unique Identifiers (UUIDs) of entities created in the repository, and also of transactions logged in the history. This identifier acts as the "UUID node", as specified by RFC 4122. |
| <b>Repository label</b>                                      | Defines a user-friendly label that indicates the purpose and context of the repository.                                                                                                                                                                                                                                                                                                                            |

---

See also [TIBCO EBX® main configuration file](#) [p 349]

## ***Defining users in the default directory***

If a custom user and roles directory is not defined in the EBX® main configuration file, the configuration assistant allows to define default users for the default user and roles directory.

An administrator user must be defined. You may optionally create a second user.

See also [Users and roles directory](#) [p 401]

## ***Validating the information entered***

Before proceeding with the installation of the repository, you can review the configuration of the repository and the information entered on the 'Configuration Summary' page. If you need to modify information, you can return to the previous pages using the configuration assistant < **Back** button.

Once you have verified the configuration, click the button **Install the repository** > to proceed with the installation.

## ***Installing the EBX® repository***

The repository installation is performed using the provided information. When the installation is complete, you are redirected to the repository login page.

## CHAPTER 61

# Deploying and registering TIBCO EBX® add-ons

### Note

Refer to the documentation of each add-on for additional installation and configuration information in conjunction with this documentation.

This chapter contains the following topics:

1. [Deploying an add-on module](#)
2. [Registering an add-on module](#)
3. [Deleting an add-on module](#)

## 61.1 Deploying an add-on module

### Note

**Each add-on bundle version is intended to run with a specific EBX® version and all its fix releases. Make sure that the EBX® and add-on bundle versions are compatible, otherwise the add-on registration will abort.**

The web application deployment descriptor for the add-on module must specify that class definitions and resources from the web application are to be loaded in preference to classes from the parent and server classloaders.

For example, on WebSphere Application Server, this can be done by setting `<context-priority-classloader>true</context-priority-classloader>` in the web-app element of the deployment descriptor.

On WebLogic, include `<prefer-web-inf-classes>True</prefer-web-inf-classes>` in `weblogic.xml`.

See the documentation on class loading of your application server for more information.

The EBX® add-on common JAR file, named `lib/ebx-addons.jar`, must be copied in the library directory shared by all web applications.

**Note**

The add-on log level can be managed in the [main configuration file](#) [p 357].

## 61.2 Registering an add-on module

To register a new EBX® add-on in the repository:

1. Navigate to the 'Administration' area.
2. Click the down-arrow in the navigation pane and select **Technical configuration > Add-ons registration**.
3. On the **Registered add-ons** page, click the + button to create a new entry.
4. Select the add-on you are registering, and enter its license key.

**Note**

If the EBX® repository is under a trial license, no license key is required for the add-on. The add-on will be subject to the same trial period as the EBX® repository itself.

The license key can be retrieved from the [TIBCO eDelivery](#) site.

5. Click on **Save**.

## 61.3 Deleting an add-on module

To delete an add-on module from the EBX® repository:

1. Navigate to the 'Administration' area.
2. Click the down-arrow in the navigation pane and select **Technical configuration > Add-ons registration**.
3. On the **Registered add-ons** page, tick the box corresponding to the add-on to be deleted.
4. In the 'Actions' menu, select 'Delete'.
5. Close and purge the Administration datasets related to the previously used add-on, as well as the including dataspace.

When an add-on is no longer deployed, a dataspace corresponding to the Administration dataset will then appear in the list of Reference children under the dataspace. When an add-on module is no longer deployed, it is thus necessary to close/delete and purge manually all data/dataspace related to the add-on.

---

# Technical administration

---

## CHAPTER 62

# Repository administration

This chapter contains the following topics:

1. [Technical architecture](#)
2. [Auto-increments](#)
3. [Repository management](#)
4. [Monitoring management](#)
5. [Dataspaces](#)

## 62.1 Technical architecture

### Overview

The main principles of the TIBCO EBX® technical architecture are the following:

- A Java process (JVM) that runs EBX® is limited to a single EBX® repository. This repository is physically persisted in a [supported relational database instance](#) [p 311], accessed through a [configured data source](#) [p 351].
- A repository cannot be shared by multiple JVMs at any given time. However, a failover architecture may be used. These aspects are detailed in the sections [Single JVM per repository](#) [p 375] and [Failover with hot-standby](#) [p 375]. Furthermore, to achieve horizontal scalability, an alternative is to deploy a [distributed data delivery \(D3\)](#) [p 426] environment.
- A single relational database instance can support multiple EBX® repositories (used by distinct JVMs). It is then required that they specify distinct table prefixes using the property `ebx.persistence.table.prefix`.

### See also

[Configuring the EBX® repository](#) [p 351]

[Supported databases](#) [p 311]

[Data source of the EBX® repository](#) [p 320]

## Rules for the database access and user privileges

### Attention

In order to guarantee the integrity of persisted master data, **it is strictly forbidden to perform direct SQL writes to the database**, except for specific use cases described in the section [SQL access to data in relational mode](#) [p 246].

It is required for the database user specified by the [configured data source](#) [p 351] to have the 'create/alter' privileges on tables, indexes and sequences. This allows for [automatic repository installation and upgrades](#) [p 377].

### See also

[SQL access to history](#) [p 252]

[Accessing a replica table using SQL](#) [p 259]

[SQL access to data in relational mode](#) [p 246]

[Data source of the EBX® repository](#) [p 320]

## Single JVM per repository

A repository cannot be shared by multiple JVMs. If such a situation was to occur, it would lead to unpredictable behavior and potentially even corruption of data in the repository.

EBX® performs checks to enforce this restriction. Before the repository becomes available, the repository must first acquire exclusive ownership of the relational database. After starting the repository, the JVM periodically checks that it still holds ownership of the repository.

These checks are performed by repeatedly tagging a technical table in the relational database. The shutdown command for the application server ensures that the tag on this technical table is removed. If the server shuts down unexpectedly, the tag may be left in the table. If this occurs, the server must wait several additional seconds upon restart to ensure that the table is not being updated by another live process.

### Attention

To avoid an additional wait period at the next start up, it is recommended to always properly shut down the application server.

## Failover with hot-standby

The exclusion mechanism described above is compatible with failover architectures, where only one server is active at any given time in an active/passive cluster. To ensure that this is the case, the main server must declare the property `ebx.repository.ownership.mode=failovermain`. The main server claims ownership of the repository database, as in the case of a single server.

A backup server can still start up, but it will not have access to the repository. It must declare the property `ebx.repository.ownership.mode=failoverstandby` to act as the backup server. Once started, the backup server is registered in the connection log. Its status can be retrieved using the Java API or through an HTTP request, as described in the section [Repository status information and logs](#) [p 376] below.

In order to activate the backup server and transfer exclusive ownership of the repository to it, a specific request must be issued by an HTTP request, or using the Java API:

- Using HTTP, the request must include the parameter `activationKeyFromStandbyMode`, and the value of this parameter must be equal to the value declared for the entry `ebx.repository.ownership.activationkey` in the EBX® main configuration file. See [Configuring failover](#) [p 362].

The format of the request URL must be:

```
http[s]://<host>[:<port>]/ebx?activationKeyFromStandbyMode={value}
```

- Using the Java API, call the method `RepositoryStatus.wakeFromStandbyAPI`.

If the main server is still up and accessing the database, the following applies: the backup server marks the ownership table in the database, requesting a clean shutdown for the main server (yet allowing any running transactions to finish). Only after the main server has returned ownership can the backup server start using the repository.

## Repository status information and logs

A log of all attempted Java process connections to the repository is available in the Administration area under '[History and logs](#) [p 249]' > 'Repository connection log'.

The status of the repository may be retrieved using the methods in the `RepositoryStatusAPI` API.

It is also possible to get the repository status information using an HTTP request that includes the parameter `repositoryInformationRequest` with one of following values:

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>state</b>            | <p>The state of the repository in terms of ownership registration.</p> <ul style="list-style-type: none"> <li>• D: Java process is stopped.</li> <li>• O: Java process has exclusive ownership of the database.</li> <li>• S: Java process is started in failover standby mode, but is not yet allowed to interact with the repository.</li> <li>• N: Java process has tried to take ownership of the database but failed because another process is holding it.</li> </ul> |
| <b>heart_beat_count</b> | <p>The number of times that the repository has made contact since associating with the database.</p>                                                                                                                                                                                                                                                                                                                                                                        |
| <b>info</b>             | <p>Detailed information for the end-user regarding the repository's registration status. The format of this information may be subject to modifications in the future without explicit warning.</p>                                                                                                                                                                                                                                                                         |



## 62.2 Auto-increments

Several technical tables can be accessed in the 'Administration' area of the EBX® user interface. These tables are for internal use only and their content should not be edited manually, unless removing obsolete or erroneous data. Among these technical tables are:

---

|                        |                                                    |
|------------------------|----------------------------------------------------|
| <b>Auto-increments</b> | Lists all auto-increment fields in the repository. |
|------------------------|----------------------------------------------------|

---

## 62.3 Repository management

### *Installation and upgrades*

#### **Automatic installation and upgrades**

By complying with the [Rules for the database access and user privileges](#) [p 375], the repository installation or upgrade is done automatically.

### *Inter-database migration*

EBX® provides a way to export the full content of a repository to another database. The export includes all dataspace, configuration datasets, and mapped tables. To operate this migration, the following guidelines must be respected:

- The source repository **must be shut down**: no EBX® server process must be accessing it; **not strictly complying with this requirement can lead to a corrupted target repository**;
- A new EBX® server process must be launched on the target repository, which must be empty. In addition to the classic Java system property `-Debx.properties`, this process must also specify `ebx.migration.source.properties`: the location of an EBX® properties file specifying the source repository. (It is allowed to provide distinct table prefixes between target and source.)
- The migration process will then take place automatically. Please note, however, that this process is not transactional: should it fail halfway, it will be necessary to delete the created objects in the target database, before starting over.
- After the migration is complete, an exception will be thrown, to force restarting the EBX® server process accessing the target repository.

Limitations:

- For technical reasons, migration to an Oracle database is only supported from another Oracle database.
- The names of the database objects representing the mapped tables (history, replication, relational) may have to be altered when migrated to the target database, to comply with the limitations of its database engine (maximum length, reserved words, ...). Such alterations will be logged during the migration process.
- As a consequence, the names specified for replicated tables in the data model will not be consistent with the adapted name in the database. The first recompilation of this data model will force to correct this inconsistency.

- Due to different representations of numeric types, values for `xs:decimal` types might get rounded if the target database engine offers a lesser precision than the source. For example, a value of `100000000.1234567890123456789` in Oracle will get rounded to `100000000.123456789012345679` in SQL Server.

## Repository backup

A global backup of the EBX® repository must be delegated to the underlying RDBMS. The database administrator must use the standard backup procedures of the underlying database.

## Archives directory

Archives are stored in a sub-directory called `archives` within the `ebx.repository.directory` (see [configuration](#) [p 349]). This directory is automatically created during the first export from EBX®.

### Attention

If manually creating this directory, make sure that the EBX® process has read-write access to it. Furthermore, the administrator is responsible for cleaning this directory, as EBX® does not maintain it.

### Note

The transfer of files between two EBX® environments must be performed using tools such as FTP or simple file copies by network sharing.

## Repository attributes

A repository has the following attributes:

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>repositoryId</b>     | Uniquely identifies a repository within the scope of the company. It is 48 bits (6 bytes) and is usually represented as 12 hexadecimal digits. This information is used for generating UUIDs (Universally Unique Identifiers) for entities created in the repository, as well as transactions logged in history tables or in the XML audit trail. This identifier acts as the 'UUID node' part, as specified by <i>RFC 4122</i> . |
| <b>repository label</b> | Provides a user-friendly label that identifies the purpose and context of the repository. For example: "Production environment".                                                                                                                                                                                                                                                                                                  |
| <b>store format</b>     | Identifies the underlying persistence system, including the current version of its structure.                                                                                                                                                                                                                                                                                                                                     |

## 62.4 Monitoring management

### ***Monitoring and cleanup of the relational database***

Some entities accumulate during the execution of EBX®.

#### **Attention**

It is the *administrator's responsibility* to monitor and clean up these entities.

### **Monitoring and reorganization**

The persistence data source of the repository must be monitored through RDBMS monitoring.

The EBX® tables specified in the default [semantic mode](#) [p 243] have their content persisted in a set of generic database tables:

- The table `${ebx.persistence.table.prefix}HOM`, in which each record represents a dataspace or a snapshot (its name is `EBX_HOM` if the property `ebx.persistence.table.prefix` is unset).
- The table `${ebx.persistence.table.prefix}BLK`, where the data of EBX® tables in semantic mode are segmented into blocks of at most 100 EBX® records (its name is `EBX_BLK` if the property `ebx.persistence.table.prefix` is unset).
- The tables `${ebx.persistence.table.prefix}HTA`, `${ebx.persistence.table.prefix}TAR` and `${ebx.persistence.table.prefix}ATB`, defining which blocks belong to a given EBX® table in a given dataspace or snapshot.

### **Database statistics**

The performance of requests executed by EBX® requires that the database has computed up-to-date statistics on its tables. Since database engines regularly schedule statistics updates, this is usually not an issue. Yet, it could be necessary to explicitly update the statistics in cases where tables are heavily modified over a short period of time (e.g. by an import creating many records).

### **Impact on UI**

Some UI components use statistics to adapt their behavior in order to prevent users from executing costly requests unwillingly.

For example, the combo box will not automatically search on user input if the table contains a large volume of records. This behavior may also occur if the database's statistics are not up to date, because a table may be considered as containing a large volume of records even if it is not actually the case.

### **Cleaning up dataspace, snapshots, and history**

A full cleanup of dataspace, snapshots, and history from the repository involves several stages:

1. Closing unused dataspace and snapshots to keep the cache to a minimal size.
2. Deleting dataspace, snapshots, and history.
3. Purging the remaining entities associated with the deleted dataspace, snapshots, and history from the repository.

## Closing unused dataspaces and snapshots

In order to keep the cache and the repository to a reasonable size, it is recommended to close any dataspaces and snapshots that are no longer required. This can be done in the following ways:

- Through the user interface, in the 'Dataspaces' area.
- From the 'Dataspaces / Snapshots' table under 'Dataspaces' in the 'Administration' area, using the **Actions** menu in the workspace. The action can be used on a filtered view of the table.
- Through the Java API, using the method `Repository.closeHomeAPI`.
- Using the data service "close dataspace" and "close snapshot" operations. See [Closing a dataspace or snapshot](#) [p 643] for more information.

Once the dataspaces and snapshots have been closed, the data can be safely removed from the repository.

### Note

Closed dataspaces and snapshots can be reopened in the 'Administration' area, under 'Dataspaces'.

## Deleting dataspaces, snapshots, and history

Dataspaces, associated history and snapshots can be permanently deleted from the repository. However, the deletion of a dataspace does not necessarily imply the deletion of its history. The two operations are independent and can be performed at different times.

### Note

The deletion of a dataspace, a snapshot, or of the history associated with them is recursive. The deletion operation will be performed on every descendant of the selected dataspace.

After the deletion of a dataspace or snapshot, some entities will remain until a repository-wide purge of obsolete data is performed. In particular, the complete history of a dataspace remains visible until a repository-wide purge is performed. Both steps, the deletion and the repository-wide purge, must be completed in order to totally remove the data and history. The process has been divided into two steps for performance issues. As the total clean-up of the repository can be time-intensive, this allows the purge execution to be initiated during off-peak periods on the server.

The process of deleting the history of a dataspace takes into account all history transactions recorded up until the deletion is submitted or until a date specified by the user. Any subsequent historized operations will not be included when the purge operation is executed. To delete new transactions, the history of the dataspace must be deleted again.

### Note

It is not possible to set a deletion date in the future. The specified date will thus be ignored and the current date will be used instead.

The deletion of dataspaces, snapshots, and history can be performed in a number of different ways:

- From the 'Dataspaces/Snapshots' table under 'Dataspaces' in the 'Administration' area, using the **Actions** menu button in the workspace. The action can be used on a filtered view of the table.

- Using the Java API, and more specifically the methods `Repository.deleteHomeAPI` and `RepositoryPurge.markHomeForHistoryPurgeAPI`.
- At the end of the data service "close dataspace" operation, using the parameters `deleteDataOnClose` and `deleteHistoryOnClose`, or at the end of a "merge dataspace" operation, using the parameters `deleteDataOnMerge` and `deleteHistoryOnMerge`.

### Purging remaining entities after a dataspace, snapshot, or history deletion

Once items have been deleted, a purge can be executed to clean up remaining data from *all* deletions performed until that point. A purge can be initiated in the following ways:

- Through the user interface, by selecting in the 'Administration' area **Actions > Execute purge** in the navigation pane.
- Using the Java API, specifically the method `RepositoryPurge.purgeAllAPI`.
- Using the task scheduler. See [Task scheduler](#) [p 415] for more information.

The purge process is logged in the directory `${ebx.repository.directory}/db.purge/`.

### Cleaning up tables having unreadable records

Some data model evolutions can lead to unreadable data:

- A column containing null values is added to the primary key of a table.
- The type of a column has changed to a different type with no possible conversion.

In these situations, records that do not match the new table definition are no longer visible, but remain persisted in the table. (This allows retrieving records if switching back to the previous table definition). When such records are encountered, an informative error is recorded in EBX® logs.

EBX® provides the option to clean the records that no longer conform to the model, once the new version of the data model is stabilized. This allows recovering space in the database and getting rid of error messages. Please proceed carefully, as this operation permanently removes all unreadable records from the selected table, and cannot be undone.

Cleaning up unreadable records is done by selecting in the 'Administration' area **Actions > Clean up unreadable records** in the navigation pane.

### Cleaning up other repository entities

It is the *administrator's responsibility* to monitor and regularly cleanup the following entities.

#### Purge

A purge can be executed to clean up the remaining data from *all* deletions, that is, deleted dataspace, snapshots and history performed up until that point. A purge can be initiated by selecting in the 'Administration' area **Actions > Execute purge** in the navigation pane.

#### Task scheduler execution reports

Task scheduler execution reports are persisted in the 'executions report' table, in the 'Task scheduler' section of the 'Administration' area. Scheduled tasks constantly add to this table as they are executed. Even when an execution terminates normally, the records are not automatically deleted. It is thus recommended to delete old records regularly.

## User interactions

User interactions are used by the EBX® component as a reliable means for an application to initiate and get the result of a service execution. They are persisted in the *ebx-interactions* administration section. It is recommended to regularly monitor the user interactions table, as well as to clean it, if needed.

## Workflow history

The workflow events are persisted in the workflow history table, in the 'Workflow' section of the 'Administration' area. Data workflows constantly add to this table as they are executed. Even when an execution terminates normally, the records are not automatically deleted. It is thus recommended to delete old records regularly.

The steps to clean history are the following

- Make sure the process executions are removed (it can be done by selecting in the 'Administration' area of Workflows **Actions > Terminate and clean this workflow** or **Actions > Clean from a date** in the navigation pane).
- Clean main processes in history (it can be done by selecting in the 'Administration' area of Workflows history **Actions > Clear from a date** or **Actions > Clean from selected workflows** in the navigation pane).
- Purge remaining entities in workflow history using 'standard EBX® purge'

See also [the standard EBX® purge](#) [p 380]

## Monitoring and clean up of file system

### Attention

In order to guarantee the correct operation of EBX®, the disk usage and disk availability of the following directories must be supervised by the administrator, as EBX® does not perform any clean up:

- **XML audit trail:** `${ebx.repository.directory}/History/`
- **Archives:** `${ebx.repository.directory}/archives/`
- **Logs:** [ebx.logs.directory](#) [p 354]
- **Temporary directory:** [ebx.temp.directory](#) [p 353]

### Attention

For **XML audit trail**, if large transactions are executed with full update details activated (default setting), the required disk space can increase.

### Attention

For pagination in the data services getChanges operation, a persistent store is used in the **Temporary directory**. Large changes may require a large amount of disk space.

See also

[\*XML audit Trail\*](#) [p 421]

[\*Tuning the EBX® repository\*](#) [p 363]

## 62.5 Dataspaces

Some dataspace administrative tasks can be performed from the 'Administration' area of EBX® by selecting 'Dataspaces'.

### ***Dataspaces/snapshots***

This table lists all the existing dataspaces and snapshots in the repository, whether open or closed. You can view and modify the information of dataspaces included in this table.

See also [\*Dataspace information\*](#) [p 94]

From this section, it is also possible to close open dataspaces, reopen previously closed dataspaces, as well as delete and purge open or closed dataspaces, associated history, and snapshots.

See also [\*Cleaning up dataspaces, snapshots, and history\*](#) [p 379]

### ***Dataspace permissions***

This table lists all the existing permission rules defined on all the dataspaces in the repository. You can view the permission rules and modify their information.

See also [\*Dataspace permissions\*](#) [p 95]

### ***Repository history***

The table 'Deleted dataspaces/snapshots' lists all the dataspaces that have already been purged from the repository.

From this section, it is also possible to delete the history of purged dataspaces.





## CHAPTER 63

# UI administration

TIBCO EBX® comes with a full user interface called [Advanced perspective](#) [p 386] that includes all available features. The interface is fully [customizable](#) [p 389] (custom logo, colors, field size, default values, etc.) and available to built-in administrators.

Access to the advanced perspective can be restricted in order to simplify the end-user experience, through [global permissions](#) [p 385], giving the possibility to grant or restrict access to functional categories. Administrators can create simplified perspectives called [recommended perspectives](#) [p 397] for end-users, containing only the features and menus they need for their daily tasks.

This chapter contains the following topics:

1. [Global permissions](#)
2. [Advanced perspective](#)
3. [Recommended perspectives](#)
4. [Custom views](#)
5. [User session management](#)

## 63.1 Global permissions

Global permission rules can be created in EBX®.

The 'Display area' property allows restricting access to areas of the user interface. To define the access rules, select 'Global permissions' in the 'Administration' area.

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Profile</b>            | Indicates on which profile the rule will be applied.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Restriction policy</b> | Indicates if the permissions defined here restrict the ones defined for other profiles. See the <a href="#">Restriction policy concept</a> [p 281] for more information.                                                                                                                                                                                                                                                |
| <b>Dataspaces</b>         | Defines permissions for the Dataspaces area.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Data Models</b>        | Defines permissions for the Data Models area.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Workflow Models</b>    | Defines permissions for the Workflow Models area.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Data Workflows</b>     | Defines permissions for the Data Workflows area.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Data Services</b>      | <p>Defines permissions for the Data Services area.</p> <p>Independently, it is also possible to:</p> <ul style="list-style-type: none"> <li>• Defines permissions for the REST built-in connector HTTP(S). This setting does not impact the REST Toolkit applications.</li> <li>• Defines permissions for the SOAP connector HTTP(S) and JMS.</li> <li>• Defines permissions for the WSDL connector HTTP(S).</li> </ul> |
| <b>Administration</b>     | Defines permissions for the Administration area.                                                                                                                                                                                                                                                                                                                                                                        |

#### Note

Permissions can be defined by administrators and by the dataspace or dataset owner.

## 63.2 Advanced perspective

The advanced perspective and its parameterization are unique. It is the parent perspective from which any [new perspective](#) [p 397] will inherit.

Children perspectives can be created from that main perspective in order to offer a customized, simplified menu to the end-users. Thanks to dataset inheritance, these simplified perspectives will receive their parameters from the advanced perspective (the root dataset). These parameters can then be overridden on the newly created simplified perspectives. Simplified perspectives can be created underneath an existing simplified perspective, thus inheriting from the parent's parameters.

See also [Inheritance](#) [p 25]

The advanced perspective is available by default to all end-users but access can be restricted.

**Note:** Administrators can always access the advanced perspective even when it is deactivated.

It is possible to configure which perspective is applied by default when users log in. This 'default perspective' is based on two criteria: 'recommended perspectives', defined by administrators and 'favorite perspectives', defined by users.

See also

[Recommended perspectives](#) [p 397]

[Favorite perspectives](#) [p 17]

## Perspective creation

To create a perspective, open the 'Select an administration feature' drop-down menu and click on the + sign to create a child dataset.

See also [Creating an inheriting child dataset](#) [p 110]

## User interface

Options are available in the Administration area for configuring the web interface, in the 'User interface' section.

### Attention

Be careful when configuring the [URL policy \(deprecated\)](#) [p 388]. If the web interface configuration is invalid, it can lead to the unusability of EBX®. If this occurs, use the "rescue mode" by setting `frontEnd.rescueMode.enable=true` in [EBX® main configuration file](#) [p 349], and accessing the following URL in your browser as a built-in administrator user: `http://.../ebx/?onwbpID=iebx-manager-rescue`.

## Session configuration

These parameters configure the user session options:

|                                      |                                                                                                                                                                    |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>User session default locale</b>   | Default session locale                                                                                                                                             |
| <b>Session time-out (in seconds)</b> | Maximum duration of user inactivity before the session is considered inactive and is terminated. A negative value indicates that the session should never timeout. |

## Interface configuration

### Entry policy

Describes the URL to access the application.

|                  |                                                                         |
|------------------|-------------------------------------------------------------------------|
| <b>Login URL</b> | If the user is not authenticated, the session is forwarded to this URL. |
|------------------|-------------------------------------------------------------------------|

The entry policy defines an EBX® login page, replacing the default one.

If defined,

- it replaces an authentication URL that may have been defined using a specific user Directory<sup>API</sup>,
- it is used to build the permalinks in the user interface,
- if the URL is full, that is, starting with `http://` or `https://`, it replaces the URL of the workflow email configuration.

### URL policy (deprecated)

Describes the URL and proxy policy. Both dynamic (servlet) and static (resources) URLs can be configured.

This configuration manner is deprecated and must be replaced by [URLs computing](#) [p 359]. After configuring the EBX® main configuration file, these configurations must be unset.

|                                        |                                                                                                                                                                                                                                                                                   |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HTTP servlet policy</b>             | Header content of the servlet HTTP request: <ul style="list-style-type: none"> <li>• if a field is not set, the default value in the environment configuration is used,</li> <li>• if a default value is not set, the value in the initial request is used.</li> </ul>            |
| <b>HTTPS servlet policy</b>            | Header content of the servlet HTTPS request: <ul style="list-style-type: none"> <li>• if a field is not set, the default value is chosen (in an environment configuration),</li> <li>• if a default value is not set, the value in the initial request is used.</li> </ul>        |
| <b>HTTP external resources policy</b>  | Header content of the external resources URL in HTTP: <ul style="list-style-type: none"> <li>• if a field is not set, the default value in the environment configuration is used,</li> <li>• if a default value is not set, the value in the initial request is used.</li> </ul>  |
| <b>HTTPS external resources policy</b> | Header content of the external resources URL in HTTPS: <ul style="list-style-type: none"> <li>• if a field is not set, the default value in the environment configuration is used,</li> <li>• if a default value is not set, the value in the initial request is used.</li> </ul> |

**Exit policy**

Describes how the application is exited.

|                                 |                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Normal redirection</b>       | Specifies the redirection URL used when exiting the session normally.                                                                          |
| <b>Error redirection</b>        | Specifies the redirection URL used when exiting the session due to an error.<br><br>This feature is now deprecated and may be ignored by EBX®. |
| <b>Redirection restrictions</b> | Specifies the list of authorized domains and whether HTTPS is mandatory for each domain.                                                       |

**Graphical interface configuration****Activation & Allowed profiles**

The 'Activated' radio button allows to activate or deactivate the perspective. When deactivated, the perspective will only be made available to the administrator.

The 'Allowed profiles' feature is used to give access to the perspective to a given profile. Several profiles can be added to the list of authorized profiles by clicking on the + icon below the numbered list.

The available perspective properties are:

|                          |                                                                                                                                        |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Activated</b>         | Indicates if the perspective is visible to authorized users.                                                                           |
| <b>Allowed profiles</b>  | The list of authorized user profiles for the perspective.                                                                              |
| <b>Allowed devices</b>   | The list of authorized devices for the perspective.<br><br>If not specified, only "EBX® Web Application" can display this perspective. |
| <b>Default selection</b> | The menu item that is selected by default.<br><br>This property is not available for the advanced perspective.                         |

**Application locking**

EBX® availability status:

|                               |                                                                                                                                 |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Availability status</b>    | This application can be closed to users during maintenance (but still remain open to administrators). Takes effect immediately. |
| <b>Unavailability message</b> | Message displayed to users when access is restricted to administrator profiles.                                                 |

**Security policy**

EBX® access security policy. These parameters only apply to new HTTP sessions.

|                               |                                                                                                                                                                                                                                                  |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IP access restriction</b>  | Restricts access to designated IP addresses (see IP pattern below).                                                                                                                                                                              |
| <b>IP restriction pattern</b> | Regular expression representation of IP addresses authorized to access EBX®. For example, ((127\.\0\.\0\.\1)   (192\.\168\.\*\.\*)) grants access to the local machine and the network IP range 192.168.*.*.                                     |
| <b>Unique session control</b> | Specifies whether EBX® should control the uniqueness of user sessions. When set to 'Yes', if a user does not log out before closing the browser, it will not be possible for that user to log in again until the previous session has timed out. |

**Ergonomics and layout**

EBX® ergonomics parameters:

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Max table columns to display</b>          | According to network and browser performance, adjusts the maximum number of columns to display in a table. This property is not used when a view is applied on a table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Maximum auto-width for table columns</b>  | Defines the maximum width to which a table column can auto-size during table initialization. This is to prevent columns from being too wide, which could occur for very long values, such as URLs. Users will still be able to manually resize columns beyond this value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Max expanded elements for a hierarchy</b> | Defines the maximum number of elements that can be expanded in a hierarchy when using the action "Expand all". A value less than or equal to '0' disables this parameter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default table filter</b>                  | Defines the default table filter to display in the filters list in tabular views. If modified, users must log out and log in for the new value to take effect.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Display the message box automatically</b> | Defines the message severity threshold for displaying the messages pop-up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>IE compatibility mode</b>                 | <p>Defines whether or not to compensate for Internet Explorer 8+ displaying EBX® in compatibility mode.</p> <p>In order to prevent Internet Explorer browsers from using compatibility mode when displaying the repository user interface, the meta-tag <code>http-equiv="X-UA-Compatible" content="IE=EmulateIE8"</code> is added to the header of pages. However, in some local environments, this setting may conflict with existing policies, in which case this header must be omitted by setting the parameter to 'No'. The default value is 'Yes'.</p> <p>See <a href="#">Specifying Document Compatibility Modes</a>  for more information.</p> |
| <b>Forms: width of labels</b>                | The width of labels in forms.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Forms: width of inputs</b>                | The width of form input fields in forms.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Forms: height of text areas</b>           | The height of text entry fields in forms.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

|                                                                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forms: aggregated lists</b>                                                  | The number of hidden candidate lines to be generated, available to create new instances in the list.                                                                                                                                                                                                                                |
| <b>Forms: width of HTML editor</b>                                              | The width of HTML editors in forms.                                                                                                                                                                                                                                                                                                 |
| <b>Forms: height of HTML editor</b>                                             | The height of HTML editors in forms.                                                                                                                                                                                                                                                                                                |
| <b>Searchable list selection page size</b>                                      | Maximum number of rows downloaded at each request of the searchable list selection (used for selecting foreign keys, enumerations, etc.).                                                                                                                                                                                           |
| <b>Record form: rendering mode for nodes</b>                                    | Specifies how to display non-terminal nodes in record forms. This should be chosen according to network and browser performance. For impact on page loading, link mode is light, expanded and collapsed modes are heavier. If this property is modified, users are required to log out and log in for the new value to take effect. |
| <b>Record form: display of selection and association nodes in creation mode</b> | If enabled, the selection and association nodes will be displayed in record creation forms.                                                                                                                                                                                                                                         |
| <b>Display density</b>                                                          | Defines the default display density mode for all users. If no density has been selected by the user yet, this value will be applied. Conversely, if the user already chose a density, their choice will prevail.                                                                                                                    |
| <b>Avatar displayed in the header</b>                                           | This property defines the display mode of avatars in the header. For example, it is possible to enable or disable the use of avatars in the header by updating this property. If no value is defined, the default value is 'Avatar only'. If it is a relative path, prefix it with "../" to get back to the application root URL.   |
| <b>Avatar displayed in the history</b>                                          | This property defines the display mode of avatars in the history. For example, it is possible to enable or disable the use of avatars in the history by updating this property. If no value is defined, the default value is 'Avatar only'. If it is a relative path, prefix it with "../" to get back to the application root URL. |
| <b>Avatar displayed in the workflow</b>                                         | This property defines the display mode of avatars in the workflow. For example, it is possible to enable or disable the use of avatars in the workflow by updating this property. If no value is defined, the default value is 'Avatar only'. If                                                                                    |



it is a relative path, prefix it with "../" to get back to the application root URL.

---

**Default option values**

Defines default values for options in the user interface.

**Import/Export**

---

|                                     |                                                                                                                                                                      |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CSV file encoding</b>            | Specifies the default character encoding to use for CSV file import and export.                                                                                      |
| <b>CSV : field separator</b>        | Specifies the default separator character to use for CSV file import and export.                                                                                     |
| <b>CSV : list separator</b>         | Specifies the default list separator character to use for CSV file import and export.                                                                                |
| <b>Import mode</b>                  | Specifies the default import mode.                                                                                                                                   |
| <b>Missing XML values as 'null'</b> | If 'Yes', when updating existing records, if a node is missing or empty in the imported file, the value is considered as 'null'. If 'No', the value is not modified. |

---

## Colors and themes

Customizes EBX® colors and themes.

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Web site icon URL (favicon)</b> | Sets a custom favicon. The recommended format is ICO, which is compatible with Internet Explorer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Logo URL (SVG)</b>              | Specifies the SVG image used for compatible browsers. Leave this field blank to use the PNG image, if specified. The user interface will attempt to use the specified PNG image if the browser is not SVG-compatible. If no PNG image is specified, the GIF/JPG image will be used. The logo must have a maximum height of 40px. If the height exceeds 40px, it will be cropped, not scaled, to a height of 40px. The width of the logo will determine the position of the buttons in the header. If it is a relative path, prefix it with "../" to get back to the application root URL. |
| <b>Logo URL (PNG)</b>              | Specifies the PNG image used for compatible browsers. Leave both this field and the SVG image URL field blank to use the GIF/JPG image. The user interface will use the GIF/JPG image if the browser is not PNG-compatible. The logo must have a maximum height of 40px. If the height exceeds 40px, it will be cropped, not scaled, to a height of 40px. The width of the logo will determine the position of the buttons in the header. If it is a relative path, prefix it with "../" to get back to the application root URL.                                                         |
| <b>Logo URL (GIF/JPG)</b>          | Specifies the GIF/JPG image to use when neither the PNG nor SVG are defined. The recommended formats are GIF and JPG. The logo must have a maximum height of 40px. If the height exceeds 40px, it will be cropped, not scaled, to a height of 40px. The width of the logo will determine the position of the buttons in the header. If it is a relative path, prefix it with "../" to get back to the application root URL.                                                                                                                                                               |
| <b>Main</b>                        | Main user interface theme color, used for selections and highlights.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Header</b>                      | Background color of the user interface header. By default, set to the same value as the Main color.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Workflow badge</b>              | Background and text/outline colors of new workflow task counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Primary buttons</b>             | Color of buttons selected by default. By default, set to the same value as the Main color.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

|                                           |                                                                                                                                                   |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Text of link style buttons</b>         | Text color of some buttons having a link style (the text is not dark or light, but colored). By default, set to the same value as the main color. |
| <b>Selected tab border</b>                | Border color of the selected tab. By default, set to the same value as the Main color.                                                            |
| <b>Table history view: technical data</b> | Background color of technical data cells in the table history view.                                                                               |
| <b>Table history view: creation</b>       | Background color of cells with the state 'creation' in the table history view.                                                                    |
| <b>Table history view: deletion</b>       | Background color of cells with the state 'deletion' in the table history view.                                                                    |
| <b>Table history view: update</b>         | Background color of cells with the state 'update' in the table history view.                                                                      |

### ***Child perspective menu***

An unlimited number of child perspectives can be created. Child perspectives inherit from the parameters of the 'Advanced perspective'. Some of these parameters can be overridden as detailed hereafter.

### **Activation & Allowed profiles**

See [Activation and Allowed profiles for the Advanced perspective](#) [p 389] for more information.

#### **Note**

Any specific parameter set for this perspective will override the default parameters that have been set in the 'Advanced perspective' configuration.

### **Perspective Menu**

This view displays the perspective menu. It is a hierarchical table view.

From this view, a user can create, delete or reorder menu item records.

See also [Hierarchical table view](#) [p 25]

|                          |                                                               |
|--------------------------|---------------------------------------------------------------|
| <b>Section Menu Item</b> | This is a top level menu item. It contains other menu items.  |
| <b>Menu group</b>        | This is a container for other menu items.                     |
| <b>Action Menu Item</b>  | This menu item displays a user service in the workspace area. |

## Menu item properties

When creating a record in the 'Perspective' Menu, the available perspective properties are:

|                           |                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>               | <p>The menu item type.</p> <p>See also <a href="#">Menu item types</a> [p 395]</p>                                                                                                                                                                                                                                                                                                        |
| <b>Parent</b>             | <p>The parent of the menu item.</p> <p>This property is not available for section menu items.</p>                                                                                                                                                                                                                                                                                         |
| <b>Label</b>              | <p>The menu item label.</p> <p>The label is optional for action menu items. If not specified, the label will be dynamically generated by EBX® when the menu item is displayed.</p>                                                                                                                                                                                                        |
| <b>Allowed devices</b>    | <p>The list of authorized devices for this item.</p> <p>If not specified, all devices can display this menu item. Currently only two devices are supported: "EBX® Web Application" and "EBX® GO".</p>                                                                                                                                                                                     |
| <b>Icon</b>               | <p>The icon for the menu item.</p> <p>Icon can be either "standard" (provided by EBX®) or an image, specified by a URL, that can be hosted on any web server.</p> <p>Icons size should be 16x16 pixels.</p> <p>This property is not available for section menu items.</p>                                                                                                                 |
| <b>Top separator</b>      | <p>Indicates that the menu item section has a top separator.</p> <p>This property is only available for section menu items.</p>                                                                                                                                                                                                                                                           |
| <b>Action</b>             | <p>The user service to execute when the user clicks on the menu item.</p> <p>See also <a href="#">User interface services</a> [p 565]</p> <p>If an end-user is allowed to view the perspective but not to execute the user service, an "access denied" message will be displayed when the user clicks on the menu item.</p> <p>This property is only available for action menu items.</p> |
| <b>Selection on close</b> | <p>The menu item that will be selected when the service terminates.</p> <p>Built-in services use this property when the user clicks on the 'Close' button.</p>                                                                                                                                                                                                                            |

This property is only available for action menu items.

---

## Ergonomics and layout

See [Ergonomics and layout for the Advanced perspective](#) [p 391] for more information.

### Note

Any specific parameter set for this perspective will override the default parameters that have been set in the 'Advanced perspective' configuration.

## Colors and themes

See [Colors and themes for the Advanced perspective](#) [p 394] for more information.

### Note

Any specific parameter set for this perspective will override the default parameters that have been set in the 'Advanced perspective' configuration.

## 63.3 Recommended perspectives

It is possible for a perspective administrator to configure recommended perspectives dedicated to a specific audience. These recommended perspectives are a way to choose which perspective is applied by default when a user logs in, based on their role.

However, users always have the possibility to switch between the various perspectives that are available to them and to set one as their favorite. See [Favorite perspectives](#) [p 17] for more information.

To configure recommended perspectives, go to *User interface > Recommended perspectives > Manage recommended perspectives*.

### Managing recommended perspectives

The main screen shows an ordered list of records associating a profile with a perspective. Note that the order here is important since a user can match more than one record (see [Resolution](#) [p 397] for more information).

- To add an entry, use the 'Create' action.
- To edit an entry, first select it in the list by clicking on it, then click on the 'Edit' action, or simply double-click on it.
- To remove an entry, first select it in the list, then click on the 'Delete' action.
- To move an entry, first select it in the list, then use the actions in the toolbar to the right of the list.

### Resolution

When a user logs in, the following algorithm determines which perspective is selected by default:

```
// 1) favorite perspective
IF the user has a favorite perspective
AND this perspective is active
AND the user is authorized for this perspective
  SELECT this perspective
  DONE

// 2) recommended perspective
FOR EACH association in the recommended perspectives list, in the declared order
  IF the user is in the declared profile
```

```

    AND the associated perspective is active
    AND the user is authorized for the associated perspective
    SELECT this perspective
    DONE

// 3) advanced perspective
IF the advanced perspective is active
AND the user is authorized for this perspective
    SELECT this perspective
    DONE

// 4) any perspective
SELECT any active perspective for which the user is authorized
DONE

```

## 63.4 Custom views

Users can create and manage custom views directly from the 'View' menu on tables. This administration section is the central point to manage these custom views.

### Views

This table contains all custom views defined on any table. Only a subset of fields is editable:

|                      |                                                                              |
|----------------------|------------------------------------------------------------------------------|
| <b>Documentation</b> | Localized labels and descriptions.                                           |
| <b>Owner</b>         | Defines the user(s) owning and authoring this view definition.               |
| <b>View group</b>    | Indicates the menu group in which this view is displayed in the 'View' menu. |
| <b>Share with</b>    | Defines the users allowed to select this view from their 'View' menu.        |

## Views permissions

This table allows to manage permissions relative to custom views, by data model and profile. The following permissions can be configured (the default value is applied when no permission is set for a given user):

| Permission             | Description                                                                                                                       | Default value                                                                                                                                                 |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Recommend views</b> | Allows the user to manage recommended views.                                                                                      | If the user is the dataset owner, the default value is 'Yes', otherwise it is 'No'.                                                                           |
| <b>Manage views</b>    | Defines the views the user can modify and delete.                                                                                 | If the user is a built-in administrator, the default value is 'Owned + shared', otherwise it is 'Owned'.                                                      |
| <b>Share views</b>     | Defines the views for which the user can edit the 'Share with' field.                                                             | If the user is a built-in administrator, the default value is 'Owned + shared', else if the user is the dataset owner, it is 'Owned', otherwise it is 'None'. |
| <b>Publish views</b>   | Allows the user to publish views to make them available to all users using Web components, workflow user tasks, or data services. | If the user is a built-in administrator, the default value is 'Yes', otherwise it is 'No'.                                                                    |

## 63.5 User session management

This tool lists all user sessions and allows terminating active sessions when necessary.

For example: it is possible to invalidate and terminate all currently open and active sessions for maintenance purposes. The access to the user interface can be temporarily closed, with an unavailability message being displayed, through [Application locking](#) [p 390]. After active sessions are terminated, users will not be able to reconnect and will see the unavailability message. The maintenance operation can then be performed.





## CHAPTER 64

# Users and roles directory

This chapter contains the following topics:

1. [Overview](#)
2. [Concepts](#)
3. [Default directory](#)
4. [Custom directory](#)

## 64.1 Overview

TIBCO EBX® uses a directory for user authentication and user role definition.

A default directory is provided and integrated into the EBX® repository; the 'Directory' administration section allows defining which users can connect and what their roles are.

It is also possible to integrate another type of enterprise directory.

**See also**

[\*Configuring the user and roles directory\*](#) [p 352]

[\*Custom directory\*](#) [p 404]

## 64.2 Concepts

In EBX®, a user can be a member of several roles, and a role can be shared by several users. Moreover, a role can be included into another role. The generic term *profile* is used to describe either a user or a role.

In addition to the directory-defined roles, EBX® provides the following *built-in roles*:

| Role                  | Definition                                                                                                                                                                                       |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Profile.ADMINISTRATOR | Built-in Administrator role. Allows performing general administrative tasks.                                                                                                                     |
| Profile.READ_ONLY     | Built-in read-only role. A user associated with the read-only role can only view the EBX® repository, and has no right to perform modifications in the repository.                               |
| Profile.OWNER         | Dynamic built-in owner role. This role is checked dynamically depending on the current element. It is only activated if the user belongs to the profile defined as owner of the current element. |
| Profile.EVERYONE      | All users belong to this role.                                                                                                                                                                   |

Information related to profiles is primarily defined in the directory.

### Attention

Associations between users and the built-in roles *OWNER* and *EVERYONE* are managed automatically by EBX®, and thus must not be modified through the directory.

User permissions are managed separately from the directory. See [Permissions](#) [p 273].

### See also

[profile](#) [p 21]

[role](#) [p 22]

[user](#) [p 21]

[administrator](#) [p 22]

[user and roles directory](#) [p 22]

## Policy

These properties configure the policies of the user and roles directory, for example, whether or not users can edit their own profiles.

## Users

This table lists all the users defined in the internal directory. New users can be added from there.

## Roles

This table lists all the users defined in the internal directory. New roles can be created in this table.

## 64.3 Default directory

### ***Directory content***

The default directory is represented by the dataset 'Directory', in the 'Administration' area.

This dataset contains tables for users and roles, as well as users' roles table, roles' inclusions table and salutations table.

#### **Note**

If a role inclusion cycle is detected, the role inclusion is ignored at the permission resolution. Refresh and check the directory validation report for cycle detection.

#### **Note**

Users' roles, roles' inclusions and salutations tables are [hidden by default](#) [p 546].

Depending on the policies defined, users can modify information related to their own accounts, regardless of the permissions defined on the directory dataset.

#### **Note**

It is not possible to delete or duplicate the default directory.

### ***Password recovery procedure***

In the default directory, passwords are encrypted (by default with a SHA256-like algorithm), and stored in this state. Consequently, it is impossible to retrieve lost passwords. A new password must be generated and sent to the user.

There are two options for this procedure:

1. A notification email is sent to the administrator, the administrator manually changes the password and sends the new password to the user.
2. A procedure automatically generates a new password and sends it to the user.

By default, the first option is used. To activate the second option, specify the property `ebx.password.remind.auto=true` in the [TIBCO EBX® main configuration file](#) [p 349].

#### **Note**

For security reasons, the password recovery procedure is not available for administrator profiles. If required, use the administrator recovery procedure instead.

### ***Administrator recovery procedure***

If all the 'login/password' credentials of the administrators are lost, a special procedure must be followed. A specific directory class redefines an administrator user with login 'admin' and password 'admin'.

To activate this procedure:

- Specify the following property in the [TIBCO EBX® main configuration file](#) [p 349]:  
`ebx.directory.factory=  
com.orchestranetworks.service.directory.DirectoryDefaultRecoverFactory`

- Start EBX® and wait until the procedure completes.
- Reset the 'ebx.directory.factory' property.
- Restart EBX® and connect using the 'admin' account.

**Note**

While the 'ebx.directory.factory' property is set for the recovery procedure, authentication of users will be denied.

## 64.4 Custom directory

As an alternative to the default directory, it is possible to integrate a specific company directory. For example, an LDAP instance, a relational database or a specific directory model instantiated into EBX®.

See also *DirectoryFactory*<sup>API</sup>

## CHAPTER 65

# Data model administration

This chapter contains the following topics:

1. [Administrating publications and versions](#)
2. [Migration of previous data models in the repository](#)
3. [Schema evolutions](#)

## 65.1 Administrating publications and versions

Technical data related to data model publications and versions can be accessed in the *Administration* section by an administrator.

*Data Modeling* contains the following two tables:

- *Publications*. Stores the publications available in the repository.
- *Versions*. Stores the versions of the data models available in the repository.

These tables are read-only but it is however possible to delete manually a publication or a version.

**Important:** If a publication or a version is deleted, then the content of associated datasets will become unavailable. So this technical data must be deleted with caution.

It is possible to spread this technical data to other TIBCO EBX® repositories exporting an archive from an EBX® repository and importing it to another one. It may be useful for propagating the evolutions of data models to other repositories.

## 65.2 Migration of previous data models in the repository

In versions before 5.2.0, published data models not depending on a module were generated in the file system directory `${ebx.repository.directory}/schemas/`, with the name of the data model (*product.xsd* for example if the data model is named *Product*). Since the 5.2.0 version, this kind of data model is now fully managed within EBX® through *Publications*. That is, republishing an existing data model migrates it as a *Publication* and redirects linked datasets to the new embedded data model. The previous XML Schema Document located in `${ebx.repository.directory}/schemas/` is renamed and suffixed with *toDelete*, meaning that the document is no longer used and can be safely deleted.

## 65.3 Schema evolutions

It is crucial to evaluate the impact of data model changes on the administration side. The following points are to be considered:

### ***Impacts on data persistence***

Administration tasks can be related to the database cleanup after a modification of the models. The following links describe how the evolutions of data models are managed at the persistence level: [Cleaning up tables having unreadable records](#) [p 381] and [Purging master tables in the database](#) [p 409].

### ***Impacts on side features***

Some components rely heavily on the data models and can be impacted by their evolutions. Some examples are: the user interface, the WSDL documents, existing archives, etc.

The 'Administration' section offers the possibility to manage some of these components (such as the views), whereas other components fall out of the administrator's scope, such as archives, WSDL files, etc.

## CHAPTER 66

# Database mapping administration

This chapter contains the following topics:

1. [Overview](#)
2. [Renaming columns in the database](#)
3. [Purging columns in the database](#)
4. [Renaming master tables in the database](#)
5. [Renaming auxiliary tables in the database](#)
6. [Purging master tables in the database](#)

## 66.1 Overview

Information and services relative to database mapping can be found in the *Administration* area.

See also

[Mapped modes](#) [p 241]

*DatabaseMapping*<sup>API</sup>

## 66.2 Renaming columns in the database

This feature is available on the 'Columns' table records, under the 'Actions' menu. It allows renaming a column in the database.

The administrator can specify the name of each column of the data model in the database for mapped modes.

Once the service is selected on a record, a summary screen displays information regarding the selected column and the administrator is prompted to enter a new name for the column in the database.

### Note

It is required that the new identifier begins with a letter.

Besides, the new name must be a valid column identifier, which depends on the naming rules of the underlying RDBMS.

See also *DatabaseMapping*<sup>API</sup>

## 66.3 Purging columns in the database

This feature is available on the 'Columns' table records, under the 'Actions' menu. It allows purging columns in mapped structures.

A column can be purged if it has been disabled for mapped modes.

A column is disabled for mapped modes when:

- the corresponding field has been removed from the data model, or
- the corresponding field has been changed in the data model, in a way that is not compatible (for example: its data type has been modified), or
- the defined mapped modes have been disabled locally on the corresponding fields, using the elements `osd:history` and `osd:replication`.

See also

[\*Disabling history on a specific field or group\*](#) [p 250]

[\*Disabling replication on a specific field or group\*](#) [p 259]

Note that this behavior will change for aggregated lists:

- when deactivating a complex aggregated list, its inner fields will still be in the `LIVING` state, whereas the list node is disabled. As lists are considered as auxiliary tables in the mapping system, this information can be checked in the 'Tables' table,
- on the other hand, when the deactivation is just for inner nodes of the list, then the list will remain `LIVING`, while its children will be `DISABLED IN MODEL`.

A column can be purged only if its own state is `DISABLED IN MODEL`, or if it is an inner field of a `DISABLED IN MODEL` list.

## 66.4 Renaming master tables in the database

This feature allows renaming master tables for both relational and history modes in the database. However, it is not available for replicated tables since their names are specified in the data model.

Both features are available on the 'Tables' table records, under the 'Actions' menu.

Master tables are database tables used for persisting the tables of the data model.

The administrator can specify in the database the name of each master table corresponding to a table of the data model.



Once the service is selected on a record, a summary screen displays information regarding the selected master table and the administrator is prompted to enter a new name for the master table in the database.

#### Note

It is required that the new identifier begins with a letter and with the repository prefix.

For history tables, it is also required that the repository prefix is followed by the history tables prefix.

For relational tables, it is also required that the repository prefix is followed by the relational tables prefix.

Besides, the new name must be a valid table identifier, which depends on the naming rules of the underlying RDBMS.

## 66.5 Renaming auxiliary tables in the database

This feature allows renaming history auxiliary tables in the database. This feature is neither available for replicated tables since their names are specified in the data model, nor for the relational mode, as aggregated lists are never supported in this mode.

This feature is available on the 'Tables' table records, under the 'Actions' menu.

Auxiliary tables are database tables used for persisting aggregated lists.

The administrator can specify in the database the name of each auxiliary table corresponding to an aggregated list of the data model.

Once the service is selected on a record, a summary screen displays information regarding the selected auxiliary table and the administrator is prompted to enter a new name for the auxiliary table in the database.

#### Note

It is required that the new identifier begins with a letter.

It is required that the new identifier begins with the repository prefix.

It is also required that the repository prefix is followed by the history tables prefix.

Besides, the new name must be a valid table identifier, which depends on the naming rules of the underlying RDBMS.

## 66.6 Purging master tables in the database

This feature allows purging history and relational tables in the database if they are no longer used.

It is available on the 'Tables' table records, under the 'Actions' menu, and is only available for master tables. This feature only applies to master tables. When a master table is purged, all its auxiliary tables are purged as well.

A mapped table can be purged in the database only if it has been disabled for the corresponding mapped mode.

To disable the mapped mode for a table, follow the procedure hereafter.

For history mode:

- Deactivate historization of the table in the data model, or
- Remove the table from the data model

For relational mode:

- Remove the table from the data model, or
- Deactivate the relational mode on the data model. As data models in semantic mode cannot be used for relational dataspace, it is thus necessary to create a dataset on a semantic dataspace using this modified data model. TIBCO EBX® will then detect that the relational mode was previously used, and will therefore propose the relational table database resources for purge.

## CHAPTER 67

# Workflow management

This chapter contains the following topics:

1. [Workflows](#)
2. [Interactions](#)
3. [Workflow history](#)

## 67.1 Workflows

To define general parameters for the execution of data workflows, the management of workflow publications, or to oversee data workflows in progress, navigate to the 'Administration' area. Click on the down arrow in the navigation pane and select *Workflow management > Workflows*.

**Note**

In cases where unexpected inconsistencies arise in the workflow execution technical tables, data workflows may encounter errors. It may then be necessary to run the operation 'Clean up inconsistencies in workflow execution tables' from the 'Actions' menu in the navigation pane under *Administration > Workflow Management > Workflows*.

### **Execution of workflows**

Various tables can be used to manage the data workflows that are currently in progress. These tables are accessible in *Workflow management > Workflows* in the navigation pane.

See also [Administration of data workflows](#) [p 173]

### **Workflows table**

The 'Workflows' table contains instances of all data workflows in the repository, including those invoked as sub-workflows. A data workflow is a particular execution instance of a workflow model publication. This table provides access to the data context variables for all data workflows. It can be used to access the status of advancement of the data workflow in terms of current variable values, and in case of a data workflow suspension, to modify the variable values.

From the 'Actions' menu of the 'Workflows' table, it is possible to clear the completed data workflows that are older than a given date, by selecting the 'Clean from a date' service. This service automatically ignores the active data workflows.

## Tokens table

The 'Tokens' table allows managing the progress of data workflows. Each token marks the current step being executed in a running data workflow, as well as the current state of the data workflow.

See also [token](#) [p 29]

## Work items table

The 'Work items' table contains all the work items associated with user tasks that currently exist. If necessary, you can manually allocate a work item to a user from this table in the case of a blockage in a data workflow. It is preferable, however, to use the buttons in the workspace of the 'Data workflows' area whenever possible to allocate, reallocate, and deallocate work items.

See also [work item](#) [p 29]

## Waiting workflows table

The 'Waiting workflows' table contains all the workflows waiting for an event. If needed, a service is available to clean this table: this service deletes all lines associated with a deleted workflow.

See also [wait task](#) [p 28]

## Comment table

The 'Comments' table contains the user's comments for main workflows and their sub-workflows.

## Workflow publications

The 'Workflow publications' table is a technical table that contains all the workflow model publications of the repository. This table associates published workflow models with their snapshots. It is not recommended to directly modify this table, but rather to use the actions available in the workflow modeling area to make changes to publications.

## Configuration

### Email configuration

In order for email notifications to be sent during the data workflow execution, the following settings must be configured under 'Email configuration':

- The URL definition field is used to build links and value mail variables in the workflow.
- The 'From email' field must be completed with the email address that will be used to send email notifications.

### Interface customization

#### Modeling default values

The default value for some properties can be customized in this section.

The administrator has the possibility to define the default values to be used when a new workflow model or workflow step is created in the 'Workflow Modeling' section.

## Work items views

Specific columns are available in the inbox and in the monitoring work items tables, in the 'Data workflows' section.

10 specific columns are available. For each specific column, a customized label can be defined.

## Priorities configuration

The property 'Default priority' defines how data workflows and their work items across the repository display if they have no priority level. For example, if this property is set to the value 'Normal', any workflow and work item with no priority will appear to have the 'Normal' priority.

The 'priorities' table defines all priority levels available to data workflows in the repository. As many integer priority levels as needed can be added, along with their labels, which will appear when users hover over the priority icon in the work item tables. The icons that correspond to each priority level can also be selected, either from the set provided by TIBCO EBX®, or by specifying a URL to an icon image file.

## Temporal tasks

Under 'Temporal tasks', the polling interval for time-dependent tasks in the workflow can be set, such as deadlines and reminders. If no interval value is set, the 'in progress' steps are checked every hour.

## Workflow inbox counter configuration

The workflow inbox counter is refreshed asynchronously, even if the end-user does not launch any action. To adjust it, two parameters need to be set:

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Cache expiry (seconds)</b>                       | Expiration time (in seconds) before a new update of the inbox cache. Please note that this parameter can impact the CPU load and performance since the computation time can be costly for a repository with many work items. If no value is defined, the default value is 600.                                                                                                                                                                                                                      |
| <b>User interface refresh periodicity (seconds)</b> | Refresh time (in seconds) between two updates of the inbox counter in the user interface. Please note that this refresh concerns all inbox counters in the user interface: inbox counters of the custom perspective, header inbox counter and Data Workflows inbox counter for the advanced perspective. If no value is defined, default value is 5. If the value is zero (or negative), the refresh is disabled. Also, the modification will only be effective after a logout/login from the user. |

Also, please note that some actions can force the inbox counter to refresh:

- access on **Data workflows**
- access on any subdivision of the **Data workflows section**
- accept or reject a work item
- launch a workflow

These parameters are accessible in *Workflow management > Workflows > Configuration > Temporal tasks* in the navigation pane.

## 67.2 Interactions

To manage workflow interactions, navigate to the Administration area. Click the down arrow in the navigation pane and select the entry *Workflow management > Interactions*.

An *interaction* is generated automatically for every work item that is created. It is a local data context of a work item and is accessible from an EBX® session. When a work item is executed, the user performs the assigned actions based upon its interaction, independently of the workflow engine. User tasks define mappings for their input and output parameters to link interactions with the overall data contexts of data workflows.

Interactions can be useful for monitoring the current parameters of work items. For example, an interaction can be updated manually by a trigger or a user service.

## 67.3 Workflow history

To view the history data workflow execution, browse the 'Administration' area. Click on the down arrow in the navigation pane and select *Workflow management > Workflow history*.

The 'Workflows' table contains all actions that have been performed during the execution of workflows.

This data can be viewed graphically or textually. It is especially useful to view the states of various objects related to workflows at a given moment. This includes actions on work items, variables in the data context, as well as tokens. In case of an error, a technical log is available.

### **Clean history**

From the 'Actions' menu of the 'Workflows' table, the history of completed data workflows older than a given date can be cleared by selecting the 'Clear from a date' service.

Only the history of workflows that have been previously cleaned (e.g. their execution data deleted) is cleared. This service automatically ignores the history associated with existing workflows. It is necessary to clear data workflows before clearing the associated history, by using the dedicated service 'Clear from a date' from the 'Workflows' table. Also, a scheduled 'Clear from a date' can be used with the built-in scheduled task *SchedulerPurgeWorkflowMainHistory*.

Please note that only main processes are cleaned. In order to remove sub-processes and all related data, it will be necessary to run a 'standard EBX® purge'.

See also [How to clean workflow history](#) [p 382]

#### **Note**

An API is available to fetch the history of a workflow. Direct access to the underlying workflow history SQL tables is not supported. See **WorkflowEngine.getProcessInstanceHistory** WorkflowEngine.getProcessInstanceHistory<sup>API</sup>.

## CHAPTER 68

# Task scheduler

This chapter contains the following topics:

1. [Overview](#)
2. [Configuration from EBX®](#)
3. [Cron expression](#)
4. [Task definition](#)
5. [Task configuration](#)

## 68.1 Overview

TIBCO EBX® offers the ability to schedule programmatic tasks.

**Note**

In order to avoid conflicts and deadlocks, tasks are scheduled in a single queue.

## 68.2 Configuration from EBX®

The declaration of schedules and tasks is done by selecting 'Task scheduler' in the 'Administration' area.

- **Schedules:** defines scheduling using "cron expressions".
- **Tasks:** configures tasks, including parametrizing task instances and user profiles for their execution.
- **Scheduled tasks:** current schedule, including task scheduling activation/deactivation.
- **Execution reports:** reports of each scheduled task run that appear immediately after the task is triggered. The reports include actions to interrupt, pause, or resume running tasks, when made available by the task definition.

## 68.3 Cron expression

(An extract of the [Quartz Scheduler](#) documentation)

The task scheduler uses "cron expressions", which can create firing schedules such as: "At 8:00am every Monday through Friday" or "At 1:30am every last Friday of the month".

## Format

A cron expression is a string comprised of 6 or 7 fields separated by a white space. Fields can contain any of the allowed values, along with various combinations of the allowed special characters for that field. The fields are as follows:

| Field Name   | Mandatory | Allowed Values   | Allowed Special Characters |
|--------------|-----------|------------------|----------------------------|
| Seconds      | Yes       | 0-59             | , - * /                    |
| Minutes      | Yes       | 0-59             | , - * /                    |
| Hours        | Yes       | 0-23             | , - * /                    |
| Day of month | Yes       | 0-31             | , - * ? / L W              |
| Month        | Yes       | 1-12 or JAN-DEC  | , - * /                    |
| Day of week  | Yes       | 1-7 or SUN-SAT   | , - * ? / L #              |
| Year         | No        | empty, 1970-2099 | , - * /                    |

A cron expression can be as simple as this: "0 \* \* \* \* ?",

or more complex, like this: "0/5 14,18,3-39,52 \* ? JAN,MAR,SEP MON-FRI 2002-2010".

### Note

The legal characters and the names of months and days of the week are not case sensitive. MON is the same as mon.

## Special characters

A cron expression is a string comprised of 6 or 7 fields separated by a white space. Fields can contain any of the allowed values, along with various combinations of the allowed special characters for that field. The fields are as follows:

- \* ("all values") - used to select all values within a field. For example, "\*" in the Minutes field means "every minute".
- ? ("no specific value") - useful when you need to specify something in one of the two fields in which the character is allowed, but not the other. For example, if I want my trigger to fire on a particular day of the month (say, the 10th), but don't care what day of the week that happens to be, I would put "10" in the day-of-month field, and "?" in the day-of-week field. See the examples below for clarification.
- - - used to specify ranges. For example, "10-12" in the hour field means "the hours 10, 11 and 12".
- , - used to specify additional values. For example, "MON,WED,FRI" in the day-of-week field means "the days Monday, Wednesday, and Friday".
- / - used to specify increments. For example, "0/15" in the seconds field means "the seconds 0, 15, 30, and 45". And "5/15" in the seconds field means "the seconds 5, 20, 35, and 50". You can also



specify '/' after the "**character - in this case**" is equivalent to having '0' before the '/'. '1/3' in the day-of-month field means "fire every 3 days starting on the first day of the month".

- **L** ("last") - has different meaning in each of the two fields in which it is allowed. For example, the value "L" in the day-of-month field means "the last day of the month" - day 31 for January, day 28 for February on non-leap years. If used in the day-of-week field by itself, it simply means "7" or "SAT". But if used in the day-of-week field after another value, it means "the last xxx day of the month" - for example "6L" means "the last friday of the month". When using the 'L' option, it is important not to specify lists, or ranges of values, as you'll get confusing results.
- **W** ("weekday") - used to specify the weekday (Monday-Friday) nearest the given day. As an example, if you were to specify "15W" as the value for the day-of-month field, the meaning is: "the nearest weekday to the 15th of the month". So if the 15th is a Saturday, the trigger will fire on Friday the 14th. If the 15th is a Sunday, the trigger will fire on Monday the 16th. If the 15th is a Tuesday, then it will fire on Tuesday the 15th. However if you specify "1W" as the value for day-of-month, and the 1st is a Saturday, the trigger will fire on Monday the 3rd, as it will not 'jump' over the boundary of a month's days. The 'W' character can only be specified when the day-of-month is a single day, not a range or list of days.

#### Note

The 'L' and 'W' characters can also be combined in the day-of-month field to yield 'LW', which translates to "last weekday of the month".

- **#** - used to specify "the nth" day-of-week day of the month. For example, the value of "6#3" in the day-of-week field means "the third Friday of the month" (day 6 = Friday and "#3" = the 3rd one in the month). Other examples: "2#1" = the first Monday of the month and "4#5" = the fifth Wednesday of the month. Note that if you specify "#5" and there is not 5 of the given day-of-week in the month, then no firing will occur that month.

## Examples

| Expression               | Meaning                                                                                                                              |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| 0 0 12 * * ?             | Fire at 12pm (noon) every day.                                                                                                       |
| 0 15 10 ? * *            | Fire at 10:15am every day.                                                                                                           |
| 0 15 10 * * ?            | Fire at 10:15am every day.                                                                                                           |
| 0 15 10 * * ? *          | Fire at 10:15am every day.                                                                                                           |
| 0 15 10 * * ? 2005       | Fire at 10:15am every day during the year 2005.                                                                                      |
| 0 * 14 * * ?             | Fire every minute starting at 2pm and ending at 2:59pm, every day.                                                                   |
| 0 0/5 14 * * ?           | Fire every 5 minutes starting at 2pm and ending at 2:55pm, every day.                                                                |
| 0 0/5 14,18 * * ?        | Fire every 5 minutes starting at 2pm and ending at 2:55pm, AND fire every 5 minutes starting at 6pm and ending at 6:55pm, every day. |
| 0 0-5 14 * * ?           | Fire every minute starting at 2pm and ending at 2:05pm, every day.                                                                   |
| 0 10,44 14 ? 3 WED       | Fire at 2:10pm and at 2:44pm every Wednesday in the month of March.                                                                  |
| 0 15 10 ? * MON-FRI      | Fire at 10:15am every Monday, Tuesday, Wednesday, Thursday and Friday.                                                               |
| 0 15 10 15 * ?           | Fire at 10:15am on the 15th day of every month.                                                                                      |
| 0 15 10 L * ?            | Fire at 10:15am on the last day of every month.                                                                                      |
| 0 15 10 ? * 6L           | Fire at 10:15am on the last Friday of every month.                                                                                   |
| 0 15 10 ? * 6L 2002-2005 | Fire at 10:15am on every last friday of every month during the years 2002, 2003, 2004 and 2005.                                      |
| 0 15 10 ? * 6#3          | Fire at 10:15am on the third Friday of every month.                                                                                  |
| 0 0 12 1/5 * ?           | Fire at 12pm (noon) every 5 days every month, starting on the first day of the month.                                                |
| 0 11 11 11 11 ?          | Fire every November 11th at 11:11am.                                                                                                 |

**Note**

Pay attention to the effects of '?' and '\*' in the day-of-week and day-of-month fields.

Support for specifying both a day-of-week and a day-of-month value is not complete (you must currently use the '?' character in one of these fields).

Be careful when setting fire times between the hours of the morning when "daylight savings" changes occur in your locale (for US locales, this would typically be the hour before and after 2:00 AM - because the time shift can cause a skip or a repeat depending on whether the time moves back or jumps forward).

## 68.4 Task definition

EBX® scheduler comes with some predefined tasks.

Custom scheduled tasks can be added by the means of **scheduler** Package `com.orchestranetworks.schedulerAPI` Java API.

The declaration of schedules and tasks is done by selecting 'Task scheduler' in the 'Administration' area.

## 68.5 Task configuration

A user must be associated with a task definition; this user will be used to generate the **session** `SessionAPI` that will run the task.

**Note**

The user will not be authenticated, and no password is required. As a consequence, a user with no password set in the directory can only be used to run scheduled tasks.

A custom task can be parameterized by means of a JavaBean specification (getter and setter).

Supported parameter types are:

- `java.lang.boolean`
- `java.lang.int`
- `java.lang.Boolean`
- `java.lang.Integer`
- `java.math.BigDecimal`
- `java.lang.String`
- `java.lang.Date`
- `java.net.URI`
- `java.net.URL`

Parameter values are set in XML format.



## CHAPTER 69

# Audit trail

This chapter contains the following topics:

1. [Overview](#)
2. [Update details and disk management](#)
3. [File organization](#)

## 69.1 Overview

XML audit trail is a feature that allows logging updates to XML files. An alternative history feature is also available to record table updates in the relational database; see [History](#) [p 249].

Any persistent updates performed in the TIBCO EBX® repository are logged to an audit trail XML file. Procedure executions are also logged, even if they do not perform any updates, as procedures are always considered to be transactions. The following information is logged:

- Transaction type, such as dataset creation, record modification, record deletion, specific procedure, etc.
- Dataspace or snapshot on which the transaction is executed.
- Transaction source. If the action was initiated by EBX®, this source is described by the user identity, HTTP session identifier and client IP address. If the action was initiated programmatically, only the user's identity is logged.
- Optional "trackingInfo" value regarding the session
- Transaction date and time (in milliseconds);
- Transaction UUID (conform to the Leach-Salz variant, version 1);
- Error information; if the transaction has failed.
- Details of the updates performed. If there are updates and if history detail is activated, see next section.

## 69.2 Update details and disk management

The audit trail is able to describe all updates made in the EBX® repository, at the finest level. Thus, the XML files can be quite large and the audit trail directory must be carefully supervised. The following should be taken into account:

1. If an archive import is executed in non-interactive mode (without a change set), the audit trail does not detail the updates; it only specifies the archive that has been imported. In this case, if it is important to keep a fine trace of the import-replace, the archive itself must be preserved.
2. If an archive import is executed in interactive mode (with a change set), or if a dataspace is merged to its parent, the resulting log size will nearly triple the unzipped size of the archive. Furthermore, for consistency concerns, each transaction is logged to a temporary file (in the audit trail directory) before being moved to the main file. Therefore, EBX® requires *at least six times the unzipped size of the largest archive that may be imported*.
3. In the context of a custom procedure that performs many updates not requiring auditing, it is possible for the developer to disable the detailed history using the method `ProcedureContext.setHistoryActivationAPI`.

See also [EBX® monitoring](#) [p 382]

## 69.3 File organization

All audit trail files are stored in the directory `${ebx.repository.directory}/History`.

### "Closed" audit files

Each file is named as follows:

```
<yyyy-mm-dd>-part<nn>.xml
```

where `<yyyy-mm-dd>` is the file date and `<nn>` is the file index for the current day.

### Writing to current audit files

When an audit file is being written, the XML structure implies working in an "open mode". The XML elements of the modifications are added to a text file named:

```
<yyyy-mm-dd>-part<nn>Content.txt
```

The standard XML format is still available in an XML file that references the text file. This file is named:

```
<yyyy-mm-dd>-part<nn>Ref.xml
```

These two files are then re-aggregated in a "closed" XML file when the repository has been cleanly shut down, or if EBX® is restarted.

### Example of an audit directory

```
2004-04-05-part00.xml
2004-04-05-part01.xml
2004-04-06-part00.xml
2004-04-06-part01.xml
2004-04-06-part02.xml
2004-04-06-part03.xml
2004-04-07-part00.xml
2004-04-10-part00.xml
2004-04-11-part00Content.txt
2004-04-11-part00Ref.xml
```

## CHAPTER 70

---

# Other

This chapter contains the following topics:

1. [Lineage](#)
2. [Event broker](#)

### 70.1 Lineage

To administer lineage, three tables are accessible:

- **Authorized profiles:** Profiles must be added to this table to be used for data lineage WSDL generation.
- **History:** Lists the general data lineage WSDLs and their configuration.
- **JMS location:** Lists the JMS URL locations.

### 70.2 Event broker

#### ***Overview***

TIBCO EBX® offers the ability to receive notifications and information related to specific events using the event broker feature. This feature consists in sending notifications related to EBX® core events to the subscriber according to their chosen topics.

## Terminology

|                     |                                                                                                                                                                                                                   |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Event broker</b> | Notification component for loosely-coupled event handling. Consists of dispatching fired events from EBX® core to concerned subscribers. The event broker is mainly used for monitoring and statistical purposes. |
| <b>Topic</b>        | Corresponds to the EBX® event type that contains messages. The number of subscribers registered to a topic is unlimited.                                                                                          |
| <b>Subscriber</b>   | Client implementation in the modules that receive the events related to the subscribed topic(s).                                                                                                                  |

## Topics

|                               |                                                                                                                                                                       |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Dataspace and snapshot</b> | Corresponds to operations in the dataspace and in the snapshot, such as: create, close, reopen, delete, archive export and archive import (only for dataspace merge). |
| <b>Repository</b>             | Corresponds to operations in the repository, such as: start-up and purge.                                                                                             |
| <b>User session</b>           | Corresponds to the operations related to user authentication, such as: login and logout.                                                                              |

## Administration

The management console is located under 'Event broker' in the 'Administration' area. It contains three tables: 'Topics', 'Subscribers' and 'Subscriptions'.

All content is read-only, except for the following operations:

- Topics and subscribers can be manually activated or deactivated using dedicated services.
- Subscribers that are no longer registered to the broker can be deleted.



---

# Distributed Data Delivery (D3)

---

## CHAPTER 71

---

# Introduction to D3

This chapter contains the following topics:

1. [Overview](#)
2. [D3 terminology](#)
3. [Known limitations](#)


## 71.1 Overview

TIBCO EBX® offers the ability to send data from an EBX® instance to other instances. Using a broadcast action, it also provides an additional layer of security and control to the other features of EBX®. It is particularly suitable for situations where data governance requires the highest levels of data consistency, approvals and the ability to rollback.

### ***D3 architecture***

A typical D3 installation consists of one primary node and multiple replica nodes. In the primary node, a Data Steward declares which dataspace must be broadcast, as well as which user profile is allowed to broadcast them to the replica nodes. The Data Steward also defines delivery profiles, which are groups of one or more dataspace.

Each replica node must define from which delivery profile it receives broadcasts.



### ***Involving third-party systems***

The features of D3 also allow third-party systems to access the data managed in EBX® through data services. Essentially, when a system consumes the data of a delivery dataspace, the data is transparently redirected to the last broadcast snapshot. This ensures a more controlled and reliable view of the managed data.

Third-party systems can either access data directly through the primary node or through a replica node. Thus, a physical architecture consisting of a primary node and no replica nodes is possible.

### ***Protocols***

If JMS is activated, the conversation between a primary node and a replica node is based on SOAP over JMS, while archive transfer is based on JMS binary messages.

If JMS is not activated, conversation between a primary node and a replica node is based on SOAP over HTTP(S), while binary archive transfer is based on TCP sockets. If HTTPS is used, make sure that the target node connector is correctly configured by enabling SSL with a trusted certificate.

See also [\*JMS for distributed data delivery \(D3\)\*](#) [p 435]

## 71.2 D3 terminology

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>broadcast</b>                | Send a publication of an official snapshot of data from a primary node to replica nodes. The broadcast transparently and transactionally ensures that the data is transferred to the replica nodes.                                                                                                                                                                                                                                                                                                                      |
| <b>delivery dataspace</b>       | A delivery dataspace is a dataspace that can be broadcast to authenticated and authorized users using a dedicated action. By default, when a data service accesses a delivery dataspace on any node, it is redirected to the last snapshot that was broadcast. See <a href="#">Data services</a> [p 433].                                                                                                                                                                                                                |
| <b>delivery profile</b>         | A delivery profile is a logical name that groups one or more delivery dataspaces. Replica nodes subscribe to one or more delivery profiles.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>cluster delivery mode</b>    | Synchronization with subscribed replica nodes is performed in a two-phase commit transactional process. This delivery mode is designed to respond to a high volume of queries using load balancing and/or fault tolerance. It ensures the consistency of data in the cluster between replica nodes and their primary node delivery dataspace. Primary and replica nodes use the same last broadcast snapshots.                                                                                                           |
| <b>federation delivery mode</b> | Synchronization is performed in a single phase, and with each registered replica node independently. This delivery mode is designed to be used with geographically distributed and/or heterogeneous architectures where response time and network availability cannot be guaranteed. At any one time, replica nodes can be at different last broadcast snapshots. The synchronization processes are thus independent of one another and replay of individual replica nodes are performed for certain broadcast failures. |
| <b>Primary node</b>             | An instance of EBX® that can define one or more delivery dataspace, and to which replica nodes can subscribe. A primary node can also act as a regular EBX® server.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Replica node</b>             | An instance of EBX® attached to a primary node, in order to receive delivery dataspace broadcasts. Besides update restrictions on delivery dataspace, the replica node acts as a regular EBX® server.                                                                                                                                                                                                                                                                                                                    |

---

|                 |                                                                                                                                                               |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hub node</b> | An instance of EBX® acting as both a primary node and a replica node. Primary delivery dataspace and replica node delivery dataspace <b>must</b> be disjoint. |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|

---

## 71.3 Known limitations

### ***General limitations***

- Each replica node must have only one primary node.
- Embedded data models cannot be used in D3 dataspace. Therefore, it is not possible to create a dataset based on a publication in a D3 dataspace.
- The compatibility is not assured if at least one replica node product version is different from the primary node.

### ***Broadcast and delivery dataspace limitations***

- Access rights on dataspace are not broadcast, whereas access rights on datasets are.
- Dataspace information is not broadcast.
- Dataspace defined in relational mode cannot be broadcast.
- If a dataspace and its parent are broadcast, their parent-child relationship will be lost in the replica nodes.
- Once a snapshot has been broadcast to a replica, subsequent broadcasts of *any* snapshot with the same name will result in restoring the originally broadcast version of that same name on the replica node. That is, if the original snapshot on the primary node is purged and a new one is created with the same name and subsequently broadcast, then the content of the replica will be restored to that of the previously broadcast snapshot, and not to the latest one of the same name.
- To guarantee dataspace consistency between D3 nodes, the data model (embedded or packaged in a module) on which the broadcast contents are based, must be the same between the primary node and its replica nodes.
- On a replica delivery dataspace, if several replica nodes are registered, and if replication is enabled in data models, it will be effective for all nodes. No setting is available to activate/deactivate replication according to D3 nodes.
- Replication on replica nodes does not take part in the distributed transaction: it is automatically triggered after commit.

### ***Administration limitations***

Technical dataspace cannot be broadcast, thus the EBX® default user directory cannot be synchronized using D3.

## CHAPTER 72

# D3 broadcasts and delivery dataspace

This chapter contains the following topics:

1. [Broadcast](#)
2. [Replica node registration](#)
3. [Accessing delivery dataspace](#)

## 72.1 Broadcast

### *Scope and contents of a broadcast*

A D3 broadcast occurs at the dataspace or snapshot level. For dataspace broadcasts, D3 first creates a snapshot to capture the current state, then broadcasts this newly created snapshot.

A broadcast performs one of the following procedures depending on the situation:

- An update of the differences computed between the new broadcast snapshot and the current 'commit' one on the replica node.
- A full synchronization containing all datasets, tables, records, and permissions. This is done on the first broadcast to a given replica node, if the previous replica node commit is not known to the primary node, or on demand using the user service in '[D3] Primary node configuration'.

See also [Services on primary nodes](#) [p 449]

### *Performing a broadcast*

The broadcast can be performed:

- By the end-user, using the **Broadcast** action available in the dataspace or snapshot (this action is available only if the dataspace is registered as a delivery dataspace)
- Using custom Java code that uses `D3NodeAsMasterAPI`.

### *Conditions*

In order to be able to broadcast, the following conditions must be fulfilled:

- The authenticated user profile has permission to broadcast.

- The dataspace or snapshot to be broadcast has no validation errors.  
**Note:** Although it is not recommended, it is possible to force a broadcast of a delivery dataspace that contains validation errors. In order to do this, set the maximum severity threshold allowed in a delivery dataspace validation report under '[D3] Primary node configuration' in the 'Administration' area.
- The D3 primary node configuration has no validation errors on the following scope: the technical record of the concerned delivery dataspace and all its dependencies (dependent delivery mappings, delivery profiles and registered replica nodes).
- The dataspace or snapshot does not contain any tables in relational mode.
- There is an associated delivery profile.
- If broadcasting a dataspace, the dataspace is not locked.
- If broadcasting a snapshot, the snapshot belongs to a dataspace declared as delivery dataspace and is not already the current broadcast snapshot (though a rollback to a previously broadcast snapshot is possible).
- The dataspace or snapshot contains differences compared to the last broadcast snapshot.

## Persistence

When a primary node shuts down, all waiting or in progress broadcast requests abort, then they will be persisted on a temporary file. On startup, all aborted broadcasts are restarted.

See also [Temporary files](#) [p 451]

### Note

Broadcasts are performed asynchronously. Therefore, no information is displayed in the user interface about the success or failure of a broadcast. Nevertheless, it is possible to monitor the broadcast operations inside '[D3] Primary node configuration'. See [Supervision](#) [p 450].

## Destination

In the target replica or hub node side:

- The ebx-d3-reference dataspace identifier is the common parent of all the delivery dataspace.
- The delivery dataspace has the same identifier in primary, replica or hub nodes.
- If the delivery dataspace is missing, it will be created on the first or on the full synchronization broadcast.
- If the delivery dataspace is already existing on the first broadcast or full synchronization it will be overridden.
- If an existing dataspace with the same identifier of the delivery one is detected outside of the ebx-d3-reference, An error will be raisen.

See also [Known limitations](#) [p 430]



## 72.2 Replica node registration

### *Scope and contents*

An initialization occurs at the replica node level according to the delivery profiles registered in the TIBCO EBX® main configuration file of the replica node. When the primary node receives that initialization request, it creates or updates the replica node entry, then sends the last broadcast snapshot of all registered delivery dataspace.

#### Note

If the registered replica node repository ID or communication layer already exists, the replica node entry in the 'Registered replica nodes' technical table is updated, otherwise a new entry is created.

### *Performing an initialization*

The initialization can be done:

- Automatically at replica node server startup.
- Manually when calling the replica node service 'Register replica node'.

### *Conditions*

To be able to register, the following conditions must be fulfilled:

- The D3 mode must be 'hub' or 'slave'.
- The primary and replica node authentication parameters must correspond to the primary node administrator and replica node administrator defined in their respective directories.
- The delivery profiles defined on the replica node must exist in the primary node configuration.
- All data models contained in the registered dataspace must exist in the replica node. If embedded, the data model names must be the same. If packaged, they must be located at the same module name and the schema path in the module must be the same in both the primary and replica nodes.
- The D3 primary node configuration has no validation error on the following scope: the technical record of the registered replica node and all its dependencies (dependent delivery profiles, delivery mappings and delivery dataspace).

#### Note

To set the parameters, see the replica or hub EBX® properties in [Configuring primary, hub and replica nodes](#) [p 445].

## 72.3 Accessing delivery dataspace

### *Data services*

By default, when a data service accesses a delivery dataspace, it is redirected to the current snapshot, which is the last broadcast one. However, this default behavior can be modified either at the request level or in the global configuration.

See also [Common parameter 'disableRedirectionToLastBroadcast'](#) [p 620]

## **Access restrictions**

On the primary node, a delivery dataspace can neither be merged nor closed. Other operations are available depending on permissions. For example, modifying a delivery dataspace directly, creating a snapshot independent from a broadcast, or creating and merging a child dataspace.

On the replica node, aside from the broadcast process, no modifications of any kind can be made to a delivery dataspace, whether by the end-user, data services, or a Java program. Furthermore, any dataspace-related operations, such as merge, close, etc., are forbidden on the replica node.

## **D3 broadcast Java API**

The last broadcast snapshot may change between two calls if a broadcast has taken place in the meantime. If a fully stable view is required for several successive calls, these calls need to specifically refer to the same snapshot.

To get the last broadcast snapshot, see `D3Node.getBroadcastVersionAPI`.

## CHAPTER 73

# D3 JMS Configuration

This chapter contains the following topics:

1. [JMS for distributed data delivery \(D3\)](#)

## 73.1 JMS for distributed data delivery (D3)

To configure D3 to use JMS instead of the default HTTP and TCP protocols, you must configure the [JMS connection factory](#) [p 321] and the following queues declared in the `WEB-INF/web.xml` deployment descriptor of the 'ebx' web application.

**Note**

If the TIBCO EBX® main configuration does not activate JMS and D3 ('slave', 'hub' or 'master' node) through the properties `ebx.d3.mode`, `ebx.jms.activate` and `ebx.jms.d3.activate`, then the environment entries below will be ignored by EBX® runtime. See [JMS](#) [p 358] and [Distributed data delivery \(D3\)](#) [p 358] in the EBX® main configuration properties for more information on these properties.

## Common declarations on primary and replica nodes (for shared queues)

| Reserved resource name       | Default JNDI name                                                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| jms/EBX_D3MasterQueue        | Weblogic: EBX_D3MasterQueue<br>JBoss: java:/jms/EBX_D3MasterQueue               | D3 primary JMS queue (only for D3 mode 'slave' or 'hub'). It specifies the queue name used to send SOAP requests to the D3 primary node. The message producer sets the primary node repository ID as a value of the header field JMSType.<br>Java type: javax.jms.Queue                                                                                                                                                                                                                                            |
| jms/EBX_D3ReplyQueue         | Weblogic: EBX_D3ReplyQueue<br>JBoss: java:/jms/EBX_D3ReplyQueue                 | D3 Reply JMS queue (for all D3 modes except the 'single' mode). It specifies the name of the reply queue for receiving SOAP responses. The consumption is filtered using the header field JMSCorrelationID.<br>Java type: javax.jms.Queue                                                                                                                                                                                                                                                                          |
| jms/EBX_D3ArchiveQueue       | Weblogic: EBX_D3ArchiveQueue<br>JBoss: java:/jms/EBX_D3ArchiveQueue             | D3 JMS Archive queue (for all D3 modes except the 'single' mode). It specifies the name of the transfer archive queue used by the D3 node. The consumption is filtered using the header field JMSCorrelationID. If the archive weight is higher than the threshold specified in the property ebx.jms.d3.archiveMaxSizeInKB, the archive will be divided into several sequences. Therefore, the consumption is filtered using the header fields JMSXGroupID and JMSXGroupSeq instead.<br>Java type: javax.jms.Queue |
| jms/EBX_D3CommunicationQueue | WebLogic: EBX_D3CommunicationQueue<br>JBoss: java:/jms/EBX_D3CommunicationQueue | D3 JMS Communication queue (for all D3 modes except 'single' mode). It specifies the name of the communication queue where the requests are received. The consumption is filtered using the header field JMSType which corresponds to the current repository ID.<br>Java type: javax.jms.Queue                                                                                                                                                                                                                     |

### Note

These JNDI names are set by default, but can be modified inside the web application archive ebx.war, included in EBXForWebLogic.ear (if using Weblogic) or in EBX.ear (if using JBoss, Websphere or other application servers).

## Optional declarations on primary nodes (for replica-specific queues)

### Note

Used for ascending compatibility prior to 5.5.0 or for mono-directional queues topology.

The deployment descriptor of the primary node must be manually modified by declaring specific communication and archive queues for each replica node. It consists in adding resource names in 'web.xml' inside 'ebx.war'. The replica-specific node queues can be used by one or more replica nodes.

Resources can be freely named, but the physical names of their associated queue must correspond to the definition of replica nodes for resources `jms/EBX_D3ArchiveQueue` and `jms/EBX_D3CommunicationQueue`.


### Note

Physical queue names matching: on registration, the replica node sends the communication and archive physical queue names. These queues are matched by physical queue name among all resources declared on the primary node. If unmatched, the registration fails.


## Examples of JMS configuration

|                                           | Shared queues                                                                           | Specific queues                                                                                |
|-------------------------------------------|-----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| <b>Primary-Replica nodes architecture</b> | <a href="#">Between a primary node and two replica nodes with shared queues</a> [p 438] | <a href="#">Between a primary node and a replica node with replica-specific queues</a> [p 439] |
| <b>Hub-Hub architecture</b>               | <a href="#">Between two hub nodes with shared queues</a> [p 440]                        | <a href="#">Between two hub nodes with replica-specific queues</a> [p 441]                     |

Between a primary node and two replica nodes with shared queues



## Between a primary node and a replica node with replica-specific queues




Reserved or custom resource name (defined in 'WEB-INF/web.xml' of 'ebx.war')

**Physical name** (resource adapter / messaging)

- ← Permanent consumer filtered by repository ID (using JMS type)
- ← -- Volatile consumer (using JMS correlation Id)
- ← - - - Producer for SOAP request or archive (using queue from JNDI resource)
- ← - - - Producer for SOAP response (using queue from JMS reply to)

Between two hub nodes with shared queues




Reserved resource name (defined in 'WEB-INF/web.xml' of 'ebx.war')

- Physical name** (resource adapter / messaging)
- Permanent consumer filtered by repository ID (using JMS type)
  - Volatile consumer (using JMS correlation Id)
  - Producer for SOAP request or archive (using queue from JNDI resource)
  - Producer for SOAP response (using queue from JMS reply to)



## Between two hub nodes with replica-specific queues



Reserved or custom resource name (defined in 'WEB-INF/web.xml' of 'ebx.war')

**Physical name** (resource adapter / messaging)

- Permanent consumer filtered by repository ID (using JMS type)
- - Volatile consumer (using JMS correlation Id)
- - Producer for SOAP request or archive (using queue from JNDI resource)
- - Producer for SOAP response (using queue from JMS reply to)



## CHAPTER 74

# D3 administration

This chapter contains the following topics:

1. [Quick start](#)
2. [Configuring D3 nodes](#)
3. [Supervision](#)

## 74.1 Quick start

This section introduces the configuration of a basic D3 architecture with two TIBCO EBX® instances. Before starting, please check that each instance can work properly with its own repository.

**Note**

Deploy EBX® on two different web application containers. If both instances are running on the same host, ensure that all communication TCP ports are distinct.

### ***Declare an existing dataspace on the primary node***

The objective is to configure and broadcast an existing dataspace from a *primary* node.

This configuration is performed on the entire D3 infrastructure ([primary](#) [p 429] and [replica](#) [p 429] nodes included).

Update the `ebx.propertiesprimary` node configuration file with:

1. Define D3 mode as *primary* in key `ebx.d3.mode`.

**Note**

The *primary* node can be started after the configuration.

After authenticating as a built-in administrator, navigate within the administration tab:

1. Prerequisite: Check that the node is configured as a *primary* node (in the 'Actions' menu use 'System information' and check 'D3 mode').
2. Open the '[D3] Primary configuration' administration feature.
3. Add the dataspace to be broadcast to the 'Delivery dataspaces' table, and declare the allowed profile.
4. Add the [delivery profile](#) [p 429] to the 'Delivery profiles' table (it must correspond to a logical name) and declare the delivery mode. Possible values are: [cluster mode](#) [p 429] or [federation mode](#) [p 429].

5. Map the delivery dataspace with the delivery profile into the 'Delivery mapping' table.

#### Note

The *primary* node is now ready for the replica node(s) registration on the delivery profile.

Check that the D3 broadcast menu appears in the 'Actions' menu of the dataspace or one of its snapshots.

## Configure replica node for registration

The objective is to configure and register the *replica* node based on a delivery profile and communications settings.

Update the `ebx.properties` replica node configuration file with:

1. Define D3 mode as replica in key `ebx.d3.mode`.
2. Define the [delivery profile](#) [p 429] set on the *primary* node in key `ebx.d3.delivery.profiles` (delivery profiles must be separated by a comma and a space).
3. Define the *primary* node user authentication (must have the built-in administrator profile) for node communications in `ebx.d3.master.username` and `ebx.d3.master.password`.
4. Define [HTTP/TCP protocols](#) [p 447] for *primary* node communication, by setting a value for the property key `ebx.d3.master.url`  
(for example `http://localhost:8080/ebx-dataservices/connector`).
5. Define the *replica* node user authentication (must have the built-in administrator profile) for node communications in `ebx.d3.slave.username` and `ebx.d3.slave.password`.
6. Define [HTTP/TCP protocols](#) [p 447] for *replica* node communication, by setting a value for the property key `ebx.d3.slave.url`  
(for example `http://localhost:8090/ebx-dataservices/connector`).

#### Note

The *replica* node can be started after the configuration.

After authenticating as a built-in administrator, navigate inside the administration tab:

1. Prerequisite: Check that the node is configured as the *replica* node (in the 'Actions' menu use 'System information' and check 'D3 mode').
2. Open the '[D3] Replica configuration' administration feature.
3. Check the information on the 'Primary information' screen: No field should have the 'N/A' value.

#### Note

Please check that the model is available before broadcast (from data model assistant, it must be published).

The *replica* node is then ready for broadcast.

## 74.2 Configuring D3 nodes

### *Runtime configuration of primary and hub nodes through the user interface*

The declaration of delivery dataspace and delivery profiles is done by selecting the '[D3] Primary configuration' feature from the 'Administration' area, where you will find the following tables:

|                           |                                                                                                             |
|---------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Delivery dataspace</b> | Declarations of the dataspace that can be broadcast.                                                        |
| <b>Delivery profiles</b>  | Profiles to which replica nodes can subscribe. The delivery mode must be defined for each delivery profile. |
| <b>Delivery mapping</b>   | The association between delivery dataspace and delivery profiles.                                           |

#### Note

The tables above are read-only while some broadcasts are pending or in progress.

### *Configuring primary, hub and replica nodes*

This section details how to configure a node in its EBX® main configuration file.

See also [Overview](#) [p 349]

#### Primary node

In order to act as a *primary* node, an instance of EBX® must declare the following properties in its main configuration file.

Sample configuration for `ebx.d3.mode=master` node:

```
#####
## D3 configuration
#####
# Configuration for master, hub and slave
#####
# Optional property.
# Possibles values are single, master, hub, slave
# Default is single meaning the server will be a standalone instance.
ebx.d3.mode=master

# Optional property.
# Master cluster broadcast timeout before rollback.
# Must be a positive integer greater than or equal to 600 (10 minutes)
# and not exceeding 86400 (one day).
# Value in seconds, default is 1800 (30 minutes).
#ebx.d3.master.cluster.broadcast.timeout=1800
```

See also [primary node](#) [p 429]

#### Hub node

In order to act as a *hub* node (combination of primary and replica node configurations), an instance of EBX® must declare the following properties in its main configuration file.

### Sample configuration for `ebx.d3.mode=hub` node:

```
#####
## D3 configuration
#####
# Configuration for master, hub and slave
#####
# Optional property.
# Possibles values are single, master, hub, slave
# Default is single meaning the server will be a standalone instance.
ebx.d3.mode=hub

# Optional property.
# Master cluster broadcast timeout before rollback.
# Must be a positive integer greater than or equal to 600 (10 minutes)
# and not exceeding 86400 (one day).
# Value in seconds, default is 1800 (30 minutes).
#ebx.d3.master.cluster.broadcast.timeout=1800

#####
# Configuration dedicated to hub or slave
#####
# Profiles to subscribe to
# Mandatory property if ebx.d3.mode=hub or ebx.d3.mode=slave
ebx.d3.delivery.profiles=

# User and password to be used to communicate with the master.
# Mandatory properties if ebx.d3.mode=hub or ebx.d3.mode=slave
ebx.d3.master.username=
ebx.d3.master.password=

# User and password to be used by the master to communicate with the hub or slave.
# Mandatory property if ebx.d3.mode=hub or ebx.d3.mode=slave
ebx.d3.slave.username=
ebx.d3.slave.password=
```

See also [hub node](#) [p 430]

### Replica node

In order to act as a *replica* node, an instance of EBX® must declare the following properties in its main configuration file.

### Sample configuration for `ebx.d3.mode=slave` node:

```
#####
## D3 configuration
#####
# Configuration for master, hub and slave
#####
# Optional property.
# Possibles values are single, master, hub, slave
# Default is single meaning the server will be a standalone instance.
ebx.d3.mode=slave

#####
# Configuration dedicated to hub or slave
#####
# Profiles to subscribe to
# Mandatory property if ebx.d3.mode=hub or ebx.d3.mode=slave
ebx.d3.delivery.profiles=

# User and password to be used to communicate with the master.
# Mandatory properties if ebx.d3.mode=hub or ebx.d3.mode=slave
ebx.d3.master.username=
ebx.d3.master.password=

# User and password to be used by the master to communicate with the hub or slave.
# Mandatory property if ebx.d3.mode=hub or ebx.d3.mode=slave
ebx.d3.slave.username=
ebx.d3.slave.password=
```

See also [replica node](#) [p 429]

## Configuring the network protocol of a node

This section details how to configure the network protocol of a node in its EBX® main configuration file.

See also [Overview](#) [p 349]

### HTTP(S) and socket TCP protocols

Sample configuration for `ebx.d3.mode=hub` or `ebx.d3.mode=slave` node with HTTP(S) network protocol:

```
#####
# HTTP(S) and TCP socket configuration for D3 hub and slave
#####
# URL to access the data services connector of the master
# Mandatory property if ebx.d3.mode=hub or ebx.d3.mode=slave and JMS for D3 is not activated.
# This property will be ignored if JMS for D3 is activated.
# The URL must follow this pattern: [protocol]://[master_host]:[master_port]/ebx-dataservices/connector
# Where the possible values of 'protocol' are 'http' or 'https'.
ebx.d3.master.url=

# URL to access the data services connector of the slave
# Mandatory property if ebx.d3.mode=hub or ebx.d3.mode=slave and JMS for D3 is not activated.
# This property will be ignored if JMS for D3 is activated.
# The URL must follow this pattern: [protocol]://[slave_host]:[slave_port]/ebx-dataservices/connector
# Where the possible values of 'protocol' are 'http' or 'https'.
ebx.d3.slave.url=

# Minimum port to use to transfer archives on TCP mode.
# Must be a positive integer above zero and below 65535.
# If not set, a random port will be used.
#ebx.d3.slave.socket.range.min=

# Max port to use on TCP mode to transfer archives.
# Must be a positive integer above ebx.d3.slave.socket.range.min and below 65535.
# Mandatory if ebx.d3.slave.socket.range.min is set.
#ebx.d3.slave.socket.range.max=
```

### JMS protocol

If JMS is activated, the following properties can be defined in order to enable JMS functionalities for a D3 node.

Sample configuration for all D3 nodes with JMS network protocol:

```
#####
## JMS configuration for D3
#####
# Taken into account only if Data Services JMS is configured properly
#####
# Configuration for master, hub and slave
#####
# Default is false, activate JMS for D3
## If activated, the deployer must ensure that the entries
## 'jms/EBX_D3ReplyQueue', 'jms/EBX_D3ArchiveQueue' and 'jms/EBX_D3CommunicationQueue'
## are bound in the operational environment of the application server.
## On slave or hub mode, the entry 'jms/EBX_D3MasterQueue' must also be bound.
ebx.jms.d3.activate=false

# Change the default timeout when using reply queue.
# Must be a positive integer that does not exceed 3600000.
# Default is 10000 milliseconds.
#ebx.jms.d3.reply.timeout=10000

# Time-to-live message value expressed in milliseconds.
# This value will be set on each message header 'JMSExpiration' that defines the
# countdown before the message deletion managed by the JMS broker.
# Must be a positive integer equal to 0 or above the value of 'ebx.jms.d3.reply.timeout'.
# The value 0 means that the message does not expire.
# Default is 3600000 (one hour).
#ebx.jms.d3.expiration=3600000

# Archive maximum size in KB for the JMS body message. If exceeds, the message
# is transferred into several sequences messages in a same group, where each one does
```

```

# not exceed the maximum size defined.
# Must be a positive integer equals to 0 or above 100.
# Default is 0 that corresponds to unbounded.
#ebx.jms.d3.archiveMaxSizeInKB=

#####
# Configuration dedicated to hub or slave
#####
# Master repository ID, used to set a message filter for the concerned master when sending JMS message
# Mandatory property if ebx.jms.d3.activate=true and if ebx.d3.mode=hub or ebx.d3.mode=slave
#ebx.jms.d3.master.repositoryId=

```

See also [JMS for distributed data delivery \(D3\)](#) [p 435]



## Services on primary nodes

Services to manage a primary node are available in the 'Administration' area of the replica node under '[D3] Primary node configuration' and also in the 'Delivery dataspace' and 'Registered replica nodes' tables. The services are:

|                                               |                                                                                                                                                                                                                                                                             |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Relaunch replays</b>                       | Immediately relaunch all replays for waiting federation deliveries.                                                                                                                                                                                                         |
| <b>Delete replica node delivery dataspace</b> | <p>Delete the delivery dataspace on chosen replica nodes and/or unregister it from the configuration of the D3 primary node.</p> <p>To access the service, select a delivery dataspace from the 'Delivery dataspace' table on the primary node, then launch the wizard.</p> |
| <b>Fully resynchronize</b>                    | Broadcast the full content of the last broadcast snapshot to the registered replica nodes.                                                                                                                                                                                  |
| <b>Subscribe a replica node</b>               | Subscribe a set of selected replica nodes.                                                                                                                                                                                                                                  |
| <b>Deactivate replica nodes</b>               | <p>Remove the selected replica nodes from the broadcast scope and switch their states to 'Unavailable'.</p> <div><b>Note</b><br/>The "in progress" broadcast contexts are rolled back.</div>                                                                                |
| <b>Unregister replica nodes</b>               | <p>Disconnects the selected replica nodes from the primary node.</p> <div><b>Note</b><br/>The "in progress" broadcast contexts are rolled back.</div>                                                                                                                       |

**Note**  
The primary node services above are hidden while some broadcasts are pending or in progress.

## Services on replica nodes

Services are available in the 'Administration' area under *[D3] Configuration of replica node* to manage its subscription to the primary node and perform other actions:

|                                                                                                       |                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Register replica node</b>                                                                          | Re-subscribes the replica node to the primary node if it has been unregistered.                                                                                                                                                                                                                                                          |
| <b>Unregister replica node</b>                                                                        | Disconnects the replica node from the primary node.                                                                                                                                                                                                                                                                                      |
| <div> <b>Note</b><br/>           The "in progress" broadcast contexts are rolled back.         </div> |                                                                                                                                                                                                                                                                                                                                          |
| <b>Close and delete snapshots</b>                                                                     | Clean up a replica node delivery dataspace.<br>To access the service, select a delivery dataspace from the 'Delivery dataspace' table on the replica node, then follow the wizard to close and delete snapshots based on their creation dates.<br><b>Note:</b> The last broadcast snapshot is automatically excluded from the selection. |

## 74.3 Supervision

The last broadcast snapshot is highlighted in the snapshot table of the dataspace, it is represented by an icon displayed in the first column.

### Primary node management console

Several tables make up the management console of the primary node, located in the 'Administration' area of the primary node, under '[D3] Primary node configuration'. They are as follows:

|                                      |                                                                                                                                                                                                       |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Registered replica nodes</b>      | Replica nodes registered with the primary node. From this table, several services are available on each record.                                                                                       |
| <b>Broadcast history</b>             | History of broadcast operations that have taken place.                                                                                                                                                |
| <b>Replica node registration log</b> | History of initialization operations that have taken place.                                                                                                                                           |
| <b>Detailed history</b>              | History of archive deliveries that have taken place. The list of associated delivery archives can be accessed from the tables 'Broadcast history' and 'Initialization history' using selection nodes. |

## Primary node supervision services

Available in the 'Administration' area of the primary node under '[D3] Primary node configuration'. The services are as follows:

|                                       |                                                                                                                                           |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check replica node information</b> | Lists the replica nodes and related information, such as the replica node's state, associated delivery profiles, and delivered snapshots. |
| <b>Clear history content</b>          | Deletes all records in all history tables, such as 'Broadcast history', 'Replica node registration log' and 'Detailed history'.           |

## Replica node monitoring through the Java API

A replica node monitoring class can be created to implement actions that are triggered when the replica node's status switches to either 'Available' or 'Unavailable'. To do so, it must implement the `NodeMonitoring` interface. This class must be outside of any EBX® module and accessible from the class-loader of 'ebx.jar' and its full class name must be specified under '[D3] Replica node configuration'.

See also *NodeMonitoring<sup>API</sup>*

## Primary node notification

A D3 administrator can set up mail notifications to receive broadcast events:

- On broadcast failure,
- On federation broadcast, if replays exceed a given threshold.

The mail contains a table of events with optional links to further details.

To enable notifications, open the '[D3] Primary node configuration' dataspace from the 'Administration' area and configure the 'Notifications' group under 'Global configuration'.

The 'From email' and 'URL definition' options should also be configured by using the 'Email configuration' link.

## Log supervision

The technical supervision can be done through the log category 'ebx.d3', declared in the EBX® main configuration file. For example:

```
ebx.log4j.category.log.d3= INFO, Console, ebxFile:d3
```

See also [Configuring the EBX® logs](#) [p 354]

## Temporary files

Some temporary files, such as exchanged archives, SOAP messages, broadcast queue, (...), are created and written to the EBX® temporary directory. This location is defined in the EBX® main configuration file:

```
#####
```

```
## Directories for temporary resources.
#####
# When set, allows specifying a directory for temporary files different from java.io.tmpdir.
# Default value is java.io.tmpdir
ebx.temp.directory = \\${java.io.tmpdir}

# Allows specifying the directory containing temporary files for cache.
# If unset, the used directory is ${ebx.temp.directory}/ebx.platform.
#ebx.temp.cache.directory = ${ebx.temp.directory}/ebx.platform

# When set, allows specifying the directory containing temporary files for import.
# If unset, the used directory is ${ebx.temp.directory}/ebx.platform.
#ebx.temp.import.directory = ${ebx.temp.directory}/ebx.platform
```

---

# Security Guide

---

## CHAPTER 75

# Security Best Practices

Here is a list of best practices that are considered useful to enforce a good security level for the EBX® setup. These best practices apply to EBX® and to other environments, their configuration, protocols and policies. These are best practices in general, and may not be relevant to your particular infrastructure and security policy.

This chapter contains the following topics:

1. [Encryption algorithms](#)
2. [HTTPS](#)
3. [Installation](#)
4. [Web Server](#)
5. [Application Server](#)
6. [Java](#)
7. [Database](#)
8. [User directory and Administration rights](#)

## 75.1 Encryption algorithms

Web Server or Application Server may specify encryption algorithms when setting HTTPS parameters. Some recommendations on these algorithms are provided in section [HTTPS](#) [p 454]. Password and fields having osd:password as a type are storing hash of their value with SHA\_512 as algorithm. That includes the password of users of the default directory.

## 75.2 HTTPS

Using HTTPS for communication with clients (GUI and REST or SOAP) is recommended. All HTTP traffic should be redirected to HTTPS.

A secure [cipher suite](#) and protocols should be used whenever possible. This applies, for example, to Web Servers, Application Servers, and jdbc connections.

TLS v1.2 should be the main protocol because it's the only version that offers modern authenticated encryption (also known as AEAD).

Several obsolete cryptographic primitives must be avoided:

- Anonymous Diffie-Hellman (ADH) suites do not provide authentication,

- NULL cipher suites provide no encryption,
- Export cipher suites are insecure when negotiated in a connection, but they can also be used against a server that prefers stronger suites (the FREAK attack),
- Suites with weak ciphers (typically of 40 and 56 bits) use encryption that can easily be broken,
- RC4 is insecure,
- 3DES is slow and weak,

On the other hand, getting too restrictive on allowed cyphers may prevent some clients to connect as they may not be able to negotiate the HTTPS connection.

The following configuration is compatible with browsers supported by EBX®.

- Cipher suites: ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-POLY1305:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-SHA384:ECDHE-RSA-AES256-SHA384:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES128-SHA256
- Versions: TLSv1.2

## 75.3 Installation

Deployed components as Web Server and Application Server should be installed using a non-root or unprivileged user, and following the [principle of least privilege](#) whenever possible. For example, only necessary ports and protocols should be opened.

## 75.4 Web Server

If you have to expose web applications on the Internet, it's a good practice to protect them with a Web Server in a [demilitarized zone](#) while EBX® and the database server may be in a production zone. Here are some best practices for the configuration.

The secure cipher suite and protocols should be set according to the above section [HTTPS](#) [p 454].

It is also a best practice not to use the default configuration, and to remove any banner that may also expose the version and type of web server.

For example, on Apache2, to remove the banner (default page returned at the root), just remove the folder `/var/www/html`.

Also, on Apache2, to remove headers identifying the Web Server, the value of [ServerTokens](#) and [ServerSignature](#) from the file `security.conf` should have the following values:

```
# ServerTokens
# This directive configures what you return as the Server HTTP response
# Header. The default is 'Full' which sends information about the OS-Type
# and compiled in modules.
# Set to one of: Full | OS | Minimal | Minor | Major | Prod
# where Full conveys the most information, and Prod the least.
ServerTokens Prod

# Optionally add a line containing the server version and virtual host
# name to server-generated pages (internal error documents, FTP directory
# listings, mod_status and mod_info output etc., but not CGI generated
# documents or custom error documents).
# Set to "Email" to also include a mailto: link to the ServerAdmin.
# Set to one of: On | Off | Email
ServerSignature Off
```

The Web Server is the recommended way for setting restrictions with HTTP security headers. Be aware that headers related to the origin will impact authorized URLs for all resources returned by EBX®. That includes the content of fields of the URL type (example: image of avatar).

Here is a list of security headers and how to set them for EBX®. First, EBX® should be configured to not set any HTTP security headers. To do so, the property `ebx.security.headers.activated` must be set to 'false'.

### *X-XSS-Protection*

The `x-xss-protection` header is designed to enable the cross-site scripting (XSS) filter built into modern web browsers. Here is what the header should look like.

```
x-xss-protection: 1; mode=block
```

For version 5.9.4, if the property `ebx.security.headers.activated` is not set or set to `true`, the security header must also be unset beforehand. For version 5.9.4, if the property `ebx.security.headers.activated` is set to `false`, the security header does not need to be unset, so ignore the first line in following snippets. For previous versions, the security header must be unset beforehand.

#### Enable in Nginx

```
header always unset x-xss-protection
header always set x-xss-protection "1; mode=block"
```

#### Enable in Apache2

```
proxy_hide_header x-xss-protection;
add_header x-xss-protection "1; mode=block" always;
```

### *x-Frame-Options*

The `x-frame-options` header provides clickjacking protection by not allowing iframes to load on the site. Be aware, this may not be compatible with your configuration if EBX® is integrated through frames for example. Here is what the header should look like:

```
x-frame-options: SAMEORIGIN
```

#### Enable in Nginx

```
add_header x-frame-options "SAMEORIGIN" always;
```

#### Enable in Apache2

```
header always sets x-frame-options "SAMEORIGIN"
```

### *X-Content-Type-Options*

The `x-content-type-options` header prevents Internet Explorer and Google Chrome from sniffing a response away from the declared content-type. This helps reduce the danger of drive-by downloads and helps treat the content properly. Here is what the header looks like.

```
x-content-type-options: nosniff
```

#### Enable in Nginx

```
add_header X-Content-Type-Options "nosniff" always;
```

#### Enable in Apache2

```
header always sets X-Content-Type-Options "nosniff"
```

### *Strict-Transport-Security*

The `strict-transport-security` header is a security enhancement that restricts web browsers to access web servers solely over HTTPS. This ensures the connection cannot be established through



an insecure HTTP connection which could be vulnerable to attacks. Here is what the header should look like:

```
strict-transport-security: max-age=31536000; includeSubDomains
```

Enable in Nginx

```
add_header Strict-Transport-Security "max-age=31536000; includeSubDomains" always;
```

Enable in Apache2

```
header always sets Strict-Transport-Security "max-age=31536000; includeSubDomains"
```

*Content-Security-Policy*

The content-security-policy HTTP header provides an additional layer of security. This policy helps prevent attacks such as Cross Site Scripting (XSS) and other code injection attacks by defining content sources which are approved and thus allowing the browser to load them. Here is what the header should look like. Make sure to adapt it with your domain name (server.company.com in the example).

```
content-security-policy: default-src 'self'; font-src * data: server.company.com; img-src * data: server.company.com; script-src * 'unsafe-inline' 'unsafe-eval'; style-src * 'unsafe-inline';
```

Enable in Nginx

```
add_header Content-Security-Policy "default-src 'self'; font-src * data: server.company.com; img-src * data: server.company.com; script-src * 'unsafe-inline' 'unsafe-eval'; style-src * 'unsafe-inline';" always;
```

Enable in Apache2

```
header always sets Content-Security-Policy "default-src 'self'; font-src * data: server.company.com; img-src * data: server.company.com; script-src * 'unsafe-inline' 'unsafe-eval'; style-src * 'unsafe-inline';"
```

*Referrer-Policy*

The Referrer-Policy HTTP header governs which referrer information should be included with requests made. The Referrer-Policy tells the web browser how to handle referrer information that is sent when a user clicks on a link that leads to another page. Here is what it should look like:

```
Referrer-Policy: strict-origin
```

Enable in Nginx

```
add_header Referrer-Policy: "strict-origin" always;
```

Enable in Apache2

```
header always sets Referrer-Policy "strict-origin"
```

## 75.5 Application Server

As for Web Servers, the same best practice applies: do not expose technical information on the Application Server. For example, for Tomcat, it is recommended to fill the attribute server of connector in server.xml with a generic value as AppServer.

```
<Connector port="8080" enableLookups="false" protocol="HTTP/1.1" useBodyEncodingForURI="true" server="AppServer"/>
```

If the Application Server is exposed through HTTPS, the secure cipher suite and Protocols should be set according to the above section [HTTPS](#) [p 454].

If there is a Web Server, it is also recommended to use ports higher than 1024 and let the Web Server do proxy.

If there is no Web Server, security headers should be set by the Application Server as described above.

## 75.6 Java

It is recommended to follow the [security best practices from Oracle](#). Last supported patches should also be applied as soon as they are available especially when they include security patches. Consider using the Server JRE for server systems, such as application servers or other long-running back-end processes. The Server JRE is the same as the regular JRE except that it does not contain the web-browser plugins.

## 75.7 Database

Databases should be encrypted at rest and in transit. If there is a private key for encryption, it should not be stored in the same location as the data files. Regarding the JDBC connection, consider configuring the JDBC driver to use SSL/TLS. Contact your database administrator for detailed instructions. You should always use the last supported version of RDBMS including drivers.

## 75.8 User directory and Administration rights

For production and test platforms, EBX® must be integrated with a [custom directory](#) [p 404] to enforce the password policy of your company. The default directory can be used only for development platforms.

According to the [Separation of Duties](#) best practice, administrators can manage users and grant access but should not have any functional rights.

---

# Developer Guide

---

---

# Introduction

---

## CHAPTER 76

# Packaging TIBCO EBX® modules

An EBX® module is a standard Java EE web application, packaging various resources such as XML Schema documents, Java classes and static resources.

Since EBX® modules are web applications they benefit from features such as class-loading isolation, WAR or EAR packaging, and Web resources exposure.

This chapter contains the following topics:

1. [Module structure](#)
2. [Module declaration](#)
3. [Module registration](#)
4. [Packaged resources](#)

## 76.1 Module structure

An EBX® module contains the following files:

|                                |                                                                                                                                                                                                  |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>/WEB-INF/ebx/module.xml</b> | This mandatory document defines the main properties and services of the module. See <a href="#">Module declaration</a> [p 462].                                                                  |
| <b>/WEB-INF/web.xml</b>        | This is the standard Java EE deployment descriptor. It can perform the registration of the EBX® module when the application server is launched. See <a href="#">Module registration</a> [p 462]. |
| <b>/META-INF/MANIFEST.MF</b>   | Optional. If present, EBX® reports the 'Implementation-Title' and 'Implementation-Version' values to <i>Administration &gt; Technical configuration &gt; Modules and data models</i> .           |
| <b>/www/</b>                   | This optional directory contains all packaged resources, which are accessible via public URL. See <a href="#">Packaged resources</a> [p 464].                                                    |

Required files for Oracle WebLogic server:

**/WEB-INF/weblogic.xml**

WebLogic deployment descriptor file which activates the prefer-web-inf-classes policy, such as the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<weblogic-web-app xmlns="http://xmlns.oracle.com/weblogic/
weblogic-web-app">
  <container-descriptor>
    <prefer-web-inf-classes>true</prefer-web-inf-classes>
  </container-descriptor>
</weblogic-web-app>
```

See [weblogic.xml Deployment Descriptor Elements](#) for more information.

## 76.2 Module declaration

A module is declared using the document /WEB-INF/ebx/module.xml. For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<module xmlns="urn:ebx-schemas:module_2.4"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:ebx-schemas:module_2.4 http://schema.orchestranetworks.com/module_2.4.xsd">
  <name>moduleTest</name>
</module>
```

See the [associated schema](#) for documentation about each property. The main properties are as follows:

| Element        | Description                                                                                                                                                                                                                                                                    | Required |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| name           | Defines the unique identifier of the module in the server instance. The module name usually corresponds to the name of the web application (the name of its directory).                                                                                                        | Yes.     |
| publicPath     | Defines a path other than the module's name identifying the web application in public URLs. This path is added to the URL of external resources of the module when computing absolute URLs. If this field is not defined, the public path is the module's name, defined above. | No.      |
| services       | Declares user services using the legacy API. See <i>Declaration and configuration</i> of legacy user services. From the version 5.8.0, it is strongly advised to use the <a href="#">new user services</a> [p 565].                                                            | No.      |
| beans          | Declares reusable Java bean components. See the <a href="#">workflow package</a> [p 551].                                                                                                                                                                                      | No.      |
| ajaxComponents | Declares Ajax components. See <b>Declaring an Ajax component in a module</b> <code>UIAjaxComponent.declareInModule<sup>api</sup></code> in the Java API.                                                                                                                       | No.      |

## 76.3 Module registration

In order to be identifiable by EBX®, a module must be registered at runtime when the application server is launched. For a web application, every EBX® module must:

- contain a Java class with the annotation `@WebListener` extending the class `ModuleRegistrationListenerAPI`.

### Attention

When using the `@WebListener` annotation, ensure that the application server is configured to activate the servlet 3.0 annotation scanning for the web application. See [JSR 315: Java™ Servlet 3.0 Specification](#) for more information.

or:

- contain a Servlet extending the class `ModuleRegistrationServletAPI`;
- make a standard declaration of this servlet in the deployment descriptor `/WEB-INF/web.xml`;
- ensure that this servlet will be registered at server startup by adding the following standard element to the deployment descriptor: `<load-on-startup>1</load-on-startup>`.

Additional recommendations and information:

- The method `handleRepositoryStartup` in `ModuleRegistrationServletAPI` allows setting the logger associated with the module and defining additional behavior such as common JavaScript and CSS resources.
- The specific class extending `ModuleRegistrationServlet` must be located in the web application (under `/WEB-INF/classes` or `/WEB-INF/lib`; due to the fact that this class is internally used as a hook to the application's class-loader, to load Java classes used by the data models associated with the module).
- The application server startup process is asynchronous and web applications / EBX® modules are discovered dynamically. The EBX® repository initialization depends on this process and will wait for the registration of all used modules up to an unlimited amount of time. As a consequence, if a used module is not deployed for any reason, it must be declared in the EBX® main configuration file. For more information, see the property [Declaring modules as undeployed](#) [p 364].
- All module registrations and unregistrations are logged in the `log.kernel` category.
- If an exception occurs while loading a module, the cause is written in the application server log.
- Once the servlet is out of service, the module is unregistered and the data models and associated datasets become unavailable. Note that hot deployment/undeployment is [not supported](#) [p 309].

## Deployment descriptor example

Here is an example of a Java EE deployment descriptor (`/WEB-INF/web.xml`):

```
<?xml version="1.0" encoding="UTF-8"?>
<web-app xmlns="http://xmlns.jcp.org/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://xmlns.jcp.org/xml/ns/javaee
    http://xmlns.jcp.org/xml/ns/javaee/web-app_3_1.xsd"
  version="3.1">
  <servlet>
    <servlet-name>InitEbxServlet</servlet-name>
    <servlet-class>com.foo.RegisterServlet</servlet-class>
    <load-on-startup>1</load-on-startup>
  </servlet>
</web-app>
```

## Registration example

Here is an implementation example of the `ModuleRegistrationServlet`:

```
package com.foo;
import javax.servlet.*;
import javax.servlet.http.*;
import com.onwbp.base.repository.*;
/**
 */
public class RegisterServlet extends ModuleRegistrationServlet
{

    public void handleRepositoryStartup(ModuleContextOnRepositoryStartup aContext)
    throws OperationException
    {
        // Perform module-specific initializations here
        ...

        // Declare custom resources here
        aContext.addExternalStyleSheetResource(MyCompanyResources.COMMON_STYLESHEET_URL);
        aContext.addExternalJavaScriptResource(MyCompanyResources.COMMON_JAVASCRIPT_URL);

        aContext.addPackagedStyleSheetResource("myModule.css");
        aContext.addPackagedJavaScriptResource("myModule.js");

    }

    public void handleRepositoryShutdown()
    {
        // Release resources of the current module when the repository is shut down here
        ...
    }

    public void destroyBeforeUnregisterModule()
    {
        // Perform operations when this servlet is being taken out of service here
        ...
    }
}
```

## 76.4 Packaged resources

The packaged resources are files and documents that can be directly accessed from client browsers and can be managed and specified either as `osd:resource` fields or via the Java API. They have various types and can also be localized.

See also

*ResourceType<sup>APT</sup>*

*Type `osd:resource` [p 486]*

### Directory structure

The packaged resources must be located under the following directory structure:

1. On the first level, the directory `/www/` must be located at the root of the module (web application).
2. On the second level, the directory must specify the localization. It can be:
  - `common/` should contain all the resources to be used by default, either because they are locale-independent or as the default localization (in EBX®, the default localization is `en`, namely English);
  - `{lang}/` when localization is required for the resources located underneath, with `{lang}` to be replaced by the actual locale code; it should correspond to the locales supported by EBX®; for more information, see [Configuring EBX® localization](#) [p 353].



3. On the third level, the directory must specify the resource type. It can be:

- `jscripts/` for JavaScript resources;
- `stylesheets/` for Cascading Style Sheet (CSS) resources;
- `html/` for HTML resources;
- `icons/` for icon typed resources;
- `images/` for image typed resources.

### **Example**

In this example, the image `logowithText.jpg` is the only resource that is localized:

```
/www
├── common
│   ├── images
│   │   ├── myCompanyLogo.jpg
│   │   └── logowithText.jpg
│   ├── jscripts
│   │   └── myCompanyCommon.js
│   └── stylesheets
│       └── myCompanyCommon.css
├── de
│   └── images
│       └── logowithText.jpg
└── fr
    └── images
        └── logowithText.jpg
```



## CHAPTER 77

# Mapping to Java

This chapter contains the following topics:

1. [How to access data from Java?](#)
2. [Concurrency and isolation levels](#)
3. [Mapping of data types](#)
4. [Java bindings](#)

## 77.1 How to access data from Java?

### **Read access**

Data can be read from various generic Java classes, mainly `AdaptationAPI` and `ValueContextAPI`. The getter methods for these classes return objects that are typed according to the mapping rules described in the section [Mapping of data types](#) [p 468].

### **Write access**

Data updates must be performed in a well-managed context:

- In the context of a procedure execution, by calling the methods `setValue...` of the interface `ValueContextForUpdateAPI`, or
- During the user input validation, by calling the method `setNewValue` of the class `ValueContextForInputValidationAPI`.

### **Modification of mutable objects**

According to the mapping that is described in the [Mapping of data types](#) [p 468] section, some accessed Java objects are mutable objects. These are instances of `List`, `Date` or any `JavaBean`. Consequently, these objects can be locally modified by their own methods. However, such modifications will remain local to the returned object unless one of the above setters is invoked and the current transaction is successfully committed.

## 77.2 Concurrency and isolation levels

### ***Highest isolation level***

The highest isolation level in ANSI/ISO SQL is `SERIALIZABLE`. Three execution methods guarantee the `SERIALIZABLE` isolation level within the scope of a dataspace:

- If the client code is run inside a `ProcedureAPI` container. This is the case for every update, for exports to XML, CSV or archive, and for data services.
- If the client code accesses a dataspace that has been explicitly locked. See `LockSpecAPI`.
- If the client code accesses data in a snapshot.

#### **Note**

For custom read-only transactions that run on a dataspace, it is recommended to use `ReadOnlyProcedureAPI`.

### ***Default isolation level***

If the client code is run outside the contexts that enable `SERIALIZABLE`, its isolation level depends on the persistence mode:

- In semantic mode, the default isolation level is `READ UNCOMMITTED`.
- In relational mode, the default isolation level is the database default isolation level.

See also [Overview of modes](#) [p 243]

### ***Java access specificities***

In a Java application, a record is represented by an instance of the `Adaptation` class. This object is initially linked to the corresponding persisted record. However, unless the client code is executed in a context that enables the [SERIALIZABLE](#) [p 468] isolation level, the object can become "disconnected" from the persisted record. If this occurs and concurrent updates have been performed, they will not be reflected in the `Adaptation` object.

Therefore, it is important for the client code to either be in a `SERIALIZABLE` context, or to regularly look up or refresh the `Adaptation` object.

#### **See also**

`AdaptationHome.findAdaptationOrNullAPI`

`AdaptationTable.lookupAdaptationByPrimaryKeyAPI`

`Adaptation.getUpToDateInstanceAPI`

## 77.3 Mapping of data types

This section describes how XML Schema type definitions and element declarations are mapped to Java types.

## Simple data types

### Basic rules for simple data types

Each XML Schema simple type corresponds to a Java class, the mapping is documented in the table [XML Schema built-in simple types](#) [p 482].

**See also** `SchemaNode.createNewOccurrence`<sup>API</sup>

### Multiple cardinality on a simple element

If the attribute `maxOccurs` is greater than 1, the element is an aggregated list and the corresponding instance in Java is an instance of `java.util.List`.

Elements of the list are instances of the Java class that is determined from the mapping of the simple type (see previous section).

## Complex data types

### Complex type definitions without a class declaration

By default (no attribute `osd:class`), a terminal node of a complex type is instantiated using an internal class. This class provides a generic `JavaBean` implementation. However, if a custom client Java code has to access these values, it is recommended to use a custom `JavaBean`. To do so, use the `osd:class` declaration described in the next section.

It is also possible to transparently instantiate, read and modify the mapped Java object, with or without the attribute `osd:class`, by invoking the methods `SchemaNode.createNewOccurrence`<sup>API</sup>, `SchemaNode.executeRead`<sup>API</sup> and `SchemaNode.executeWrite`<sup>API</sup>.

### Mapping of complex types to custom JavaBeans

It is possible to map an XML Schema complex type to a custom Java class. This is done by adding the attribute `osd:class` to the complex node definition. Unless the element has `xs:maxOccurs > 1`, you must also specify the attribute `osd:access` for the node to be considered a *terminal* node. If the element has `xs:maxOccurs > 1`, it is automatically considered to be terminal.

The custom Java class must conform to the `JavaBean` protocol. This means that each child of the complex type must correspond to a `JavaBean` property of the class. Additionally, each `JavaBean` property must be a read-write property, and its implementation must ensure that the value set by the setter method is returned, as-is, by the getter method. Contextual computations are not allowed in these methods.

### Example

In this example, the Java class `com.carRental.Customer` must define the methods `getFirstName()` and `setFirstName(String)`.

A `JavaBean` can have a custom user interface within TIBCO EBX®, by using a `UIBeanEditor`<sup>API</sup>.

```
<xs:element name="customer" osd:access="RW">
  <xs:complexType name="subscriber" osd:class="com.carRental.Customer">
    <xs:sequence>
      <xs:element name="firstName" type="xs:string"/>
      ...
    </xs:sequence>
  </xs:complexType>
</xs:element>
```

## Multiple cardinality on a complex element

If the attribute `maxOccurs` is greater than 1, then the corresponding instance in Java is:

- An instance of `java.util.List` for an aggregated list, where every element in the list is an instance of the Java class determined by the [mapping of simple types](#) [p 482], or
- An instance of `AdaptationTableAPI`, if the property `osd:table` is specified.

## 77.4 Java bindings

Java bindings allow generating Java types that reflect the structure of the data model. The Java code generation can be done in the user interface. See [Generating Java bindings](#) [p 473].

### Benefits

Ensuring the link between XML Schema structure and Java code provides a number of benefits:

- **Development assistance:** Auto-completion when typing an access path to parameters, if supported by your IDE.
- **Access code verification:** All accesses to parameters are verified at code compilation.
- **Impact verification:** Each modification of the data model impacts the code compilation state.
- **Cross-referencing:** By using the reference tools of your IDE, it is easy to verify where a parameter is used.

Consequently, it is strongly recommended to use Java bindings.

### XML declaration

The specification of the Java types to be generated from the data model is included in the main schema. Each binding element defines a generation target. It must be located at, in XPath notation, `xs:schema/xs:annotation/xs:appinfo/ebxbnd:binding`, where the prefix `ebxbnd` is a reference to the namespace identified by the URI `urn:ebx-schemas:binding_1.0`. Several binding elements can be defined if you have different generation targets.

The attribute `targetDirectory` of the element `ebxbnd:binding` defines the root directory used for Java type generation. Generally, it is the directory containing the project source code, `src`. A relative path is interpreted based on the current runtime directory of the VM, as opposed to the XML schema.

See [bindings XML Schema](#).

### XML bindings example

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:ebxbnd="urn:ebx-schemas:binding_1.0">
  <xs:annotation>
    <xs:appinfo>
      <!-- The bindings define how this schema will be represented in Java.
      Several <binding> elements may be defined, one for each target. -->
      <ebxbnd:binding
        targetDirectory="._/ebx-demos/src-creditOnlineStruts-1.0/">
        <jpathConstants typeName="com.creditonline.RulesPaths">
          <nodes root="/rules" prefix="" />
        </jpathConstants>
        <jpathConstants typeName="com.creditonline.StylesheetConstants">
          <nodes root="/stylesheet" prefix="" />
        </jpathConstants>
      </ebxbnd:binding>
    </xs:appinfo>
  </xs:annotation>
  ...
```

```
</xs:schema>
```

Java constants can be defined for XML schema paths. To do so, generate one or more interfaces from a schema node, including the root node `/`. The example generates two Java path constant interfaces, one from the node `/rules` and the other from the node `/stylesheet` in the schema. Interface names are described by the element `javaPathConstants` with the attribute `typeName`. The associated node is described by the element nodes with the attribute `root`.





## CHAPTER 78

# Tools for Java developers

TIBCO EBX® provides Java developers with tools to facilitate use of the EBX® API, as well as integration with development environments.

This chapter contains the following topics:

1. [Activating the development tools](#)
2. [Data model refresh tool](#)
3. [Generating Java bindings](#)
4. [Path to a node](#)
5. [Web component link generator](#)

## 78.1 Activating the development tools

To activate the development tools, run EBX® in *development mode*. This is specified in the EBX® main configuration file [EBX® run mode](#) [p 366] using the property `backend.mode=development`.

## 78.2 Data model refresh tool

When editing the data model directly as an XML Schema document without using the data-modeling tool provided by EBX®, you can refresh it without restarting the application server.

In the 'Administration' area, select **Select > Technical configuration > Development tools > Refresh updated data models** (or **Refresh all data models**).

### Attention

Since the operation is critical regarding data consistency, refreshing the data models acquires a global exclusive lock on the repository. This means that most other operations (data access and update, validation, etc.) will wait until the completion of the data model refresh.

## 78.3 Generating Java bindings

The Java types specified by Java bindings can be generated from a dataset or a data model, by selecting **Actions > Generate Java** in the navigation pane.

See also [Java bindings](#) [p 470]

## 78.4 Path to a node

The field 'Data path' is displayed in the documentation pane of a node. This field indicates the path to the node, which can be useful when writing XPath formulas.

### Note

This field is always available to administrators.

## 78.5 Web component link generator

The 'Web component link generator' service is a user interface designed to create HTTP requests that call EBX® web components. To launch this service, select **Actions > Web component link generator** in the navigation pane.

## CHAPTER 79

# Terminology changes

A new TIBCO EBX® release can introduce new vocabulary for users. To preserve the backward compatibility, these terminology changes do not usually impact the API. Consequently, Java class names, method names, data services operation names, etc. still use the older version terminology. This chapter purpose is to facilitate the correspondence of the old term in the API to the new terms.

See also [Glossary](#) [p 21]

This chapter contains the following topics:

1. [Terminology changes in version 5.9](#)
2. [Terminology changes in version 5.0](#)

## 79.1 Terminology changes in version 5.9

| New term        | Term prior to version 5.9.0 |
|-----------------|-----------------------------|
| D3 primary node | D3 master node              |
| D3 replica node | D3 slave node               |

## 79.2 Terminology changes in version 5.0

The following table summarizes the mappings between the version 5.0.0 terminology and previous terminology:

| New term                       | Term prior to version 5.0.0 |
|--------------------------------|-----------------------------|
| <b>Dataset</b>                 | Adaptation instance         |
| <b>Child dataset</b>           | Child adaptation instance   |
| <b>Data model</b>              | Data model                  |
| <b>Dataspace</b>               | Branch                      |
| <b>Snapshot</b>                | Version                     |
| <b>Dataspace or snapshot</b>   | Home                        |
| <b>Data Workflow</b>           | Workflow instance           |
| <b>Workflow model</b>          | Workflow definition         |
| <b>Workflow publication</b>    | Workflow                    |
| <b>Data services</b>           | Data services               |
| <b>Field</b>                   | Attribute                   |
| <b>Inherited field</b>         | Inherited attribute         |
| <b>Record</b>                  | Record/occurrence           |
| <b>Validation rule</b>         | Constraint                  |
| <b>Simple/advanced control</b> | Simple/advanced constraint  |

---

# Data model

---

## CHAPTER 80

# Introduction

A data model is a structural definition of the data to be managed in the TIBCO EBX® repository. Data models contribute to EBX®'s ability to guarantee the highest level of data consistency and to facilitate data management.

Specifically, the data model is a document that conforms to the XML Schema standard (W3C recommendation). Its main features are as follows:

- A rich library of well-defined [simple data types](#) [p 481], such as integer, boolean, decimal, date, time;
- The ability to define additional [simple types](#) [p 483] and [complex types](#) [p 483];
- The ability to define simple lists of items, called [aggregated lists](#) [p 492];
- [Validation constraints](#) [p 515] (facets), for example: enumerations, uniqueness constraints, minimum/maximum boundaries.

EBX® also uses the extensibility features of XML Schema for other useful information, such as:

- [Predefined types](#) [p 484], for example: locale, resource, html;
- Definition of [tables](#) [p 495] and [foreign key constraints](#) [p 500];
- Mapping data in EBX® to Java beans;
- [Advanced validation constraints](#) [p 515] (extended facets), such as dynamic enumerations;
- Extensive [presentation information](#) [p 533], such as labels, descriptions, and error messages.

### Note

EBX® supports a subset of the W3C recommendations, as some features are not relevant to Master Data Management.

This chapter contains the following topics:

1. [Editing the data model](#)
2. [References](#)
3. [Relationship between datasets and data models](#)
4. [Pre-requisite for XML Schemas](#)
5. [Conventions](#)
6. [Schemas with reserved names](#)

## 80.1 Editing the data model

There are two different ways to define a data model:

- The data model can be defined using an XML Schema editor or through the data model assistant. The data model assistant has the advantage of being integrated into the EBX® user interface, abstracting the verbose underlying XML. For more information, see [Introduction to data models](#) [p 32]. The data model assistant allows using features that are not documented to be used outside of the DMA; e.g. Toolbars and Widgets.
- By using an external XML Schema document editor.

## 80.2 References

For an introduction to XML Schema, see the W3Schools [XML Schema Tutorial](#).

See also

[XML Schema Part 0: Primer](#)

[XML Schema Part 1: Structures](#)

[XML Schema Part 2: Datatypes](#)

## 80.3 Relationship between datasets and data models

Each root dataset is associated with a single data model. At the dataspace creation, an associated data model is selected, on which to base the dataset.

See also [Creating a dataset](#) [p 109]

## 80.4 Pre-requisite for XML Schemas

In order for an XML Schema to be accepted by EBX®, it must include a global element declaration that includes the attribute `osd:access="- -"`.

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:osd="urn:ebx-schemas:common_1.0" xmlns:fmt="urn:ebx-schemas:format_1.0">
  <xs:import namespace="urn:ebx-schemas:common_1.0"
    schemaLocation="http://schema.orchestranetworks.com/common_1.0.xsd"/>
  <xs:element name="root" osd:access="- -"/>
  ...
</xs:element>
</xs:schema>
```

## 80.5 Conventions

By convention, namespaces are always defined as follows:

| Prefix | Namespace                          |
|--------|------------------------------------|
| xs:    | http://www.w3.org/2001/XMLSchema   |
| osd:   | urn:ebx-schemas:common_1.0         |
| fmt:   | urn:ebx-schemas:format_1.0         |
| usd:   | urn:ebx-schemas:userServices_1.0   |
| emd:   | urn:ebx-schemas:entityMappings_1.0 |

## 80.6 Schemas with reserved names

Several data models in EBX® have reserved names.

All references to other data models (using the attribute `schemaLocation` for an import, include or redefine) that end with one of the following strings are reserved:

- `common_1.0.xsd`
- `org_1.0.xsd`
- `coreModel_1.0.xsd`
- `session_1.0.xsd`
- `userServices_1.0.xsd`
- `entityMappings_1.0.xsd`

These XSD files correspond to the schemas provided for the module `ebx-root-1.0`, at the path `/WEB-INF/ebx/schemas`. The attribute `schemaLocation` can reference the files at this location or a copy, if the file names are identical. This is useful if you want to avoid a module dependency on `ebx-root-1.0`.

For security reasons, EBX® uses an internal definition for these schemas to prevent any modification.



---

## CHAPTER 81

# Data types

This chapter details the data types supported by TIBCO EBX®.

See also [Tables and relationships](#) [p 495]

This chapter contains the following topics:

1. [XML Schema built-in simple types](#)
2. [XML Schema named simple types](#)
3. [XML Schema complex types](#)
4. [Extended simple types defined by EBX®](#)
5. [Complex types defined by EBX®](#)
6. [Aggregated lists](#)
7. [Including external data models](#)

## 81.1 XML Schema built-in simple types

The table below lists all the simple types defined in XML Schema that are supported by EBX®, along with their corresponding Java types.

| XML Schema type                                                      | Java class                        | Notes                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">xs:string</a>                                            | <code>java.lang.String</code>     |                                                                                                                                                                                                                                                                                             |
| <a href="#">xs:boolean</a>                                           | <code>java.lang.Boolean</code>    |                                                                                                                                                                                                                                                                                             |
| <a href="#">xs:decimal</a>                                           | <code>java.math.BigDecimal</code> | A <code>totalDigits</code> facet with a value equal to 15 is added by default to decimal fields that are contained in a mapped table (relational, historized or replicated table). However, this facet can be overwritten with a greater value in the data model.                           |
| <a href="#">xs:dateTime</a>                                          | <code>java.util.Date</code>       |                                                                                                                                                                                                                                                                                             |
| <a href="#">xs:time</a>                                              | <code>java.util.Date</code>       | The date portion of the returned Date is always set to '1970/01/01'.                                                                                                                                                                                                                        |
| <a href="#">xs:date</a>                                              | <code>java.util.Date</code>       | The time portion of the returned Date is always the beginning of the day, that is, '00:00:00'.                                                                                                                                                                                              |
| <a href="#">xs:anyURI</a>                                            | <code>java.net.URI</code>         |                                                                                                                                                                                                                                                                                             |
| <a href="#">xs:Name</a> ( <a href="#">xs:string</a> restriction)     | <code>java.lang.String</code>     |                                                                                                                                                                                                                                                                                             |
| <a href="#">xs:int</a> ( <a href="#">xs:decimal</a> restriction)     | <code>java.lang.Integer</code>    |                                                                                                                                                                                                                                                                                             |
| <a href="#">xs:integer</a> ( <a href="#">xs:decimal</a> restriction) | <code>java.lang.Integer</code>    | This mapping does not comply with the XML Schema recommendation. Although the XML Schema specification states that <code>xs:integer</code> has no value space limitation, this value space is, in fact, restricted by the Java specifications of the <code>java.lang.Integer</code> object. |

The mapping between XML Schema types and Java types are detailed in the section [Mapping of data types](#) [p 468].

## 81.2 XML Schema named simple types

Named simple types can be defined when designing a data model for redefining an existing built-in simple type. A *named simple type* can be reused in the data model.

Restrictions:

- In the data model, only the element restriction is allowed in a named simple type, and even then, only derivation by restriction is supported. Notably, the elements `list` and `union` are not supported.
- Facet definition is not cumulative. That is, if an element and its named type both define the same kind of facet, then the facet defined in the type is overridden by the local facet definition. However, this restriction does not apply to programmatic facets defined by the element `osd:constraint`. For `osd:constraint`, if an element and its named type both define a programmatic facet with different Java classes, the definition of these facets will be cumulative. Contrary to the XML Schema Specification, EBX® is not strict regarding the definition of a facet of the same kind in an element and its named type. That is, the value of a same kind of facet defined in an element is not checked according to the one defined in the named type. However, in the case of static enumerations defined both in an element and its type, the local enumeration will be replaced by the intersection between these enumerations.
- It is not possible to define different types of enumerations on both an element and its named type. For instance, you cannot specify a static enumeration in an element and a dynamic enumeration in its named type.
- It is not possible to simultaneously define a pattern facet in both an element and its named type.

## 81.3 XML Schema complex types

Complex types can be defined when designing a data model. A *named complex type* can be reused in the data model.

Restrictions:

- In the data model, only the element sequence is allowed. Notably, attribute definition is not supported.
- Type extensions are not supported in the current version of EBX®.

## 81.4 Extended simple types defined by EBX®

EBX® provides pre-defined simple data types:

| XML Schema type                                           | Java class       |
|-----------------------------------------------------------|------------------|
| osd:text ( <a href="#">xs:string</a> restriction)         | java.lang.String |
| osd:html ( <a href="#">xs:string</a> restriction)         | java.lang.String |
| osd:email ( <a href="#">xs:string</a> restriction)        | java.lang.String |
| osd:password ( <a href="#">xs:string</a> restriction)     | java.lang.String |
| osd:color ( <a href="#">xs:string</a> restriction)        | java.lang.String |
| osd:resource ( <a href="#">xs:anyURI</a> restriction)     | internal class   |
| osd:locale ( <a href="#">xs:string</a> restriction)       | java.util.Locale |
| osd:dataspaceKey ( <a href="#">xs:string</a> restriction) | java.lang.String |
| osd:datasetName ( <a href="#">xs:string</a> restriction)  | java.lang.String |

The above types are defined by the internal schema `common-1.0.xsd`. They are defined as follows:

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>osd:text</b>     | <p>This type represents textual information. For a basic <code>xs:string</code>, its default user interface in EBX® consists of a dedicated editor with several lines for input and display.</p> <pre>&lt;xs:simpleType name="text"&gt;   &lt;xs:restriction base="xs:string" /&gt; &lt;/xs:simpleType&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>osd:html</b>     | <p>This represents a character string with HTML formatting. A WYSIWYG editor is provided in EBX®.</p> <pre>&lt;xs:simpleType name="html"&gt;   &lt;xs:restriction base="xs:string" /&gt; &lt;/xs:simpleType&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>osd:email</b>    | <p>This represents an email address as specified by the <a href="#">RFC822</a> standard.</p> <pre>&lt;xs:simpleType name="email"&gt;   &lt;xs:restriction base="xs:string" /&gt; &lt;/xs:simpleType&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>osd:password</b> | <p>This represents a hashed or encrypted password. A specific editor is provided in EBX®.</p> <pre>&lt;xs:element name="password" type="osd:password" /&gt;</pre> <p>The default editor performs a hash computation using the SHA-512 algorithm. This encryption function is also available from a Java client using the method <code>DirectoryDefault.encryptString<sup>API</sup></code>.</p> <p>It is also possible for the default editor to use a different encryption mechanism by specifying a class that implements the interface <code>Encryption<sup>API</sup></code>.</p> <pre>&lt;xs:element name="password" type="osd:password"&gt;   &lt;xs:annotation&gt;     &lt;xs:appinfo&gt;       &lt;osd:uiBean class="com.orchestranetworks.ui.UIPassword"&gt;         &lt;encryptionClass&gt;package.EncryptionClassName&lt;/encryptionClass&gt;       &lt;/osd:uiBean&gt;     &lt;/xs:appinfo&gt;   &lt;/xs:annotation&gt; &lt;/xs:element&gt;</pre> <p>It is possible to specify some salt by referencing a path to another field, and by using a class the implements the interface <code>HashComputation<sup>API</sup></code>.</p> <pre>&lt;xs:element name="password" type="osd:password"&gt;   &lt;xs:annotation&gt;     &lt;xs:appinfo&gt;       &lt;osd:uiBean class="com.orchestranetworks.ui.UIPassword"&gt;         &lt;encryptionClass&gt;package.HashClassName&lt;/encryptionClass&gt;         &lt;saltPath&gt;../login&lt;/saltPath&gt;       &lt;/osd:uiBean&gt;     &lt;/xs:appinfo&gt;   &lt;/xs:annotation&gt; &lt;/xs:element&gt;</pre> |

**osd:locale**

This represents a geographical, political or cultural location. The locale type is translated into Java by the class `java.util.Locale`.

```
<xs:simpleType name="locale">
  <xs:restriction base="xs:string">
    <xs:enumeration value="ar" osd:label="Arabic" />
    <xs:enumeration value="ar_AE" osd:label="Arabic (United Arab Emirates)" />
    <xs:enumeration value="ar_BH" osd:label="Arabic (Bahrain)" />
    <xs:enumeration value="ar_DZ" osd:label="Arabic (Algeria)" />
    <xs:enumeration value="ar_EG" osd:label="Arabic (Egypt)" />
    <xs:enumeration value="ar_IQ" osd:label="Arabic (Iraq)" />
    ...
    <xs:enumeration value="vi_VN" osd:label="Vietnamese (Vietnam)" />
  >
  <xs:enumeration value="zh" osd:label="Chinese" />
  <xs:enumeration value="zh_CN" osd:label="Chinese (China)" />
  <xs:enumeration value="zh_HK" osd:label="Chinese (Hong Kong)" />
  <xs:enumeration value="zh_TW" osd:label="Chinese (Taiwan)" />
</xs:restriction>
</xs:simpleType>
```

**osd:color**

This represents a character string with hexadecimal RGB color formatting. A color picker `UIComponent` is provided in EBX®.

```
<xs:simpleType name="color">
  <xs:restriction base="xs:string" />
</xs:simpleType>
```

**osd:resource**

This represents a resource packaged in a module. For more information, see [Packaged resources](#) [p 464]. This type requires the definition of the facet [FacetOResource](#) [p 520].

```
<xs:simpleType name="resource">
  <xs:restriction base="xs:anyURI" />
</xs:simpleType>
```

**osd:dataspaceKey**

This type represents a reference to a dataspace.

```
<xs:element name="dataspaceField" type="osd:dataspaceKey" />
```

A specific editor is provided in EBX® that displays the dataspaces that can be referenced.

It is possible to specify the dataspaces that can be referenced using the element `osd:dataspaceSet` under `xs:annotation/xs:appInfo`. If the element `osd:dataspaceSet` is not defined, then by default, only open branches can be referenced.

```
<xs:element name="dataspaceField" type="osd:dataspaceKey">
  <xs:annotation>
    <xs:appinfo>
      <osd:dataspaceSet>
        <include>
          <pattern>a pattern</pattern>
          <type>all | branch | version</type>
          <includeDescendants>none | allDescendants |
allBranchDescendants | allSnapshotDescendants | branchChildren |
snapshotChildren</includeDescendants>
        </include>
        <exclude>
          <pattern>a pattern</pattern>
          <type>all | branch | version</type>
        </exclude>
      </osd:dataspaceSet>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

```

<includeDescendants>none | allDescendants |
allBranchDescendants | allSnapshotDescendants | branchChildren |
snapshotChildren</includeDescendants>
</include>
<filter osd:class="com.foo.MyDataspaceFilter">
  <param1>...</param1>
  <param2>...</param2>
</filter>
</osd:dataspaceSet>
</xs:appinfo>
</xs:annotation>
</xs:element>

```

- includes

Specifies the dataspace that can be referenced by this field. An `include` must at least be defined.

**pattern:** Specifies a pattern that filters dataspace. The pattern is checked against the name of the dataspace. This property is mandatory.

**type:** Specifies the type of dataspace that can be referenced by this field. If not defined, this restriction is applied to branches. If `all` then branches and snapshots are included. If `branch` then only branches are included. If `snapshot` then only snapshots are included. If not set, this property is `branch` by default.

**includeDescendants:** Specifies if children or descendants of the dataspace that match the specified pattern are included in the set. If `none` then neither children nor descendants of the dataspace that match the specified pattern are included. If `allDescendants` then all descendants of the dataspace that match the specified pattern are included. If `allBranchDescendants` then all descendant branches of the dataspace that match the specified pattern are included. If `allSnapshotDescendants` then all descendant snapshots of the dataspace that match the specified pattern are included. If `directBranchChildren` then only direct branches of the dataspace that match the specified pattern are included. If `directSnapshotChildren` then only direct snapshots of the dataspace that match the specified pattern are included. If not set, this property is `none` by default.

- excludes

Specifies the dataspace that cannot be referenced by this field. Excludes are ignored if no includes are defined.

**pattern:** Specifies a pattern that filters dataspace. The pattern is checked against the name of the dataspace. This property is mandatory.

**type:** Specifies the type of dataspace that can be referenced by this field. If not defined, this restriction is applied to branches. If `all` then branches and

snapshots are excluded. If `branch` then only branches are excluded. If `snapshot` then only snapshots are excluded. If not set, this property is `branch` by default.

**includeDescendants:** Specifies if children or descendants of the datasets that match the specified pattern are excluded from the set. If `none` then neither children nor descendants of the dataspace that match the specified pattern are excluded. If `allDescendants` then all descendants of the dataspace that match the specified pattern are excluded. If `allBranchDescendants` then all descendant branches of the dataspace that match the specified pattern are excluded. If `allSnapshotDescendants` then all descendant snapshots of the dataspace that match the specified pattern are excluded. If `directBranchChildren` then only direct branches of the dataspace that match the specified pattern are excluded. If `directSnapshotChildren` then only direct snapshots of the dataspace that match the specified pattern are excluded. If not set, this property is `none` by default.

- `filter`

Specifies a filter to accept or reject dataspace in the context of a dataset or a record. This filter is only used in the dedicated input component that is associated to this field. That is, this filter is not used when validating this field. A specific constraint can be used to perform specific controls on this field.

The attribute `osd:class` specifies a Java bean that implements the interface `DataspaceSetFilter`<sup>API</sup>.

It is also possible to customize validation messages and the control policy associated with this type using the element `validation` under `xs:annotation/xs:appInfo/osd:dataspaceSet`. See [Facet validation message with severity](#) [p 536] and [Control policy](#) [p 524] for more information.

---

#### **`osd:datasetName`**

This type represents a reference to a dataset.

```
<xs:element name="dataset" type="osd:datasetName" />
```

A specific editor provided in EBX® displays the datasets that can be referenced.

It is also possible to specify the datasets that can be referenced using the element `osd:datasetSet` under `xs:annotation/xs:appInfo`:

```
<xs:element name="datasetField" type="osd:datasetName">
  <xs:annotation>
    <xs:appinfo>
      <osd:datasetSet>
        <branch>productsBranch</branch>
      </osd:datasetSet>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```



```

<version>productsVersion</version>
<dataspaceSelector>../dataspaceField</dataspaceSelector>
<pattern>a pattern</pattern>
<filter osd:class="com.foo.MyDatasetFilter">
  <param1>...</param1>
  <param2>...</param2>
</filter>
</osd:datasetSet>
</xs:appinfo>
</xs:annotation>
</xs:element>

```

- **branch**

Specifies the source branch. Only datasets contained in this branch will be able to be selected by a field of the type Dataset identifier (osd:datasetName).

- **version**

Specifies the source snapshot. Only datasets contained in this snapshot will be able to be selected by a field of the type Dataset identifier (osd:datasetName).

- **dataspaceSelector**

Specifies a field in the same data model that defines the dataspace containing the datasets that can be referenced. The specified field must be of type xs:string or osd:dataspaceKey. The value of this field must comply with the representation of a persistent identifier of a dataspace or snapshot. See `HomeKey.formatAPI` for more information.

The referred node must respect the restrictions existing for dynamic facets, see [Dynamic constraints](#) [p 518].

- **includes**

Specifies the datasets that can be referenced by this field.

**pattern:** Specifies a pattern that filters datasets. The pattern is checked against the name of the datasets. This property is mandatory.

**includeDescendants:** Specifies if children or descendants of the datasets that match the specified pattern are included in the set. If none then neither children nor descendants of the datasets that match the specified pattern are excluded. If `directChildren` then only direct children of the datasets that match the specified pattern are excluded. If `allDescendants` then all descendants of the datasets that match the specified pattern are excluded. If not set, this property is none by default.

- **excludes**

Specifies the datasets that cannot be referenced by this field. Excludes are ignored if no includes are defined.

**pattern:** Specifies a pattern that filters datasets. The pattern is checked against the name of the datasets. This property is mandatory.

**includeDescendants:** Specifies if children or descendants of the datasets that match the specified pattern are included in the set. If none then neither children nor descendants of the datasets that match the specified pattern are excluded. If `directChildren` then only direct children of the datasets that match the specified pattern are excluded. If `allDescendants` then all descendants of the datasets that match the specified pattern are excluded. If not set, this property is none by default.

- **filter**

Specifies a filter to accept or reject datasets in the context of a dataset or record. This filter is only used in the dedicated input component that is associated to this field. That is, this filter is not used when validating this field. A specific constraint can be used to perform specific controls on this field.

The attribute `osd:class` specifies a Java bean that implements the interface `DatasetSetFilter`<sup>API</sup>. A validation message is added to the associated field if an input dataspace reference does not match this filter.

One of the elements `branch`, `version` or `dataspaceSelector` must be defined.

It is also possible to customize validation messages and the control policy associated with this type using the element `validation` under `xs:annotation/xs:appInfo/osd:datasetSet`. See [Facet validation message with severity](#) [p 536] and [Control policy](#) [p 524] for more information.

---

# 81.5 Complex types defined by EBX®

EBX® provides pre-defined complex data types:

| XML Schema type | Description                                                                                                                                                                                  |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| osd:UDA         | <b>User Defined Attribute:</b> This type allows any user, according to their access rights, to define a value associated with an attribute defined in a dictionary called a UDA Catalog.     |
| osd:UDACatalog  | <b>Catalog of User Defined Attributes:</b> This type consists of a table in which attributes can be specified. This catalog is used by all osd:UDA elements declared in the same data model. |

**osd:UDA**

A User Defined Attribute (UDA) supports both the minOccurs and maxOccurs attributes, as well as the attribute osd:UDACatalogPath, which specifies the path of the corresponding catalog.

```
<xs:element name="firstUDA" type="osd:UDA" minOccurs="0"
maxOccurs="unbounded" osd:UDACatalogPath="//insuranceCatalog" />
<xs:element name="secondUDA" type="osd:UDA" minOccurs="1"
maxOccurs="1"
osd:UDACatalogPath="/root/userCatalog" />
<xs:element name="thirdUDA" type="osd:UDA" minOccurs="0"
maxOccurs="1"
osd:UDACatalogPath="//userCatalog" />
```

In the manager, when working with a UDA, the editor will adapt itself to the type of the selected attribute.

**osd:UDACatalog**

Internally, a catalog is represented as a table. The parameters minOccurs and maxOccurs must be specified.

Several catalogs can be defined in the same data model.

```
<xs:element name="insuranceCatalog" type="osd:UDACatalog"
minOccurs="0" maxOccurs="unbounded">
  <xs:annotation>
    <xs:documentation xml:lang="en-US">Insurance Catalog.</
xs:documentation>
    <xs:documentation xml:lang="fr-FR">Catalog assurance.</
xs:documentation>
  </xs:annotation>
</xs:element>
<xs:element name="userCatalog" type="osd:UDACatalog" minOccurs="0"
maxOccurs="unbounded">
  <xs:annotation>
    <xs:documentation xml:lang="en-US">User catalog.</
xs:documentation>
    <xs:documentation xml:lang="fr-FR">Catalogue utilisateur.</
xs:documentation>
  </xs:annotation>
</xs:element>
```

Only the following types are available for creating new attributes:

- xs:string
- xs:boolean
- xs:decimal

- xs:dateTime
- xs:time
- xs:date
- xs:anyURI
- xs:Name
- xs:int
- osd:html
- osd:email
- osd:password
- osd:locale
- osd:text

---

## ***Restrictions on User Defined Attributes and Catalogs***

The following features are unsupported on UDA elements:

- Facets
- Functions using the osd:function property
- UI bean editors using the osd:uiBean property
- The osd:checkNullInput property
- History features
- Replication
- Inheritance features, using the osd:inheritance property

As UDA catalogs are internally considered to be tables, the restrictions that apply to tables also exist for UDACatalog elements.

## **81.6 Aggregated lists**

In XML Schema, the maximum number of times an element can occur is determined by the value of the maxOccurs attribute in its declaration. If this value is strictly greater than 1 or is unbounded, the data can have multiple occurrences. If no osd:table declaration is included, this element is called an *aggregated list*. In Java, it is then represented as an instance of the class `java.util.List`.

The following is an example of an aggregated list that defines the pricing of a loan product, depending on the amount borrowed.

```
<xs:element name="pricing" minOccurs="0" maxOccurs="unbounded"
  osd:access="RW">
  <xs:annotation>
    <xs:documentation>
      <osd:label>Pricing</osd:label>
      <osd:description>Pricing grid </osd:description>
    </xs:documentation>
  </xs:annotation>
  <xs:complexType>
    <xs:sequence>
      <xs:element name="amount" type="xs:int">
        <xs:annotation>
          <xs:documentation>
            <osd:label>Amount borrowed</osd:label>
          </xs:documentation>
        </xs:annotation>
      </xs:element>
    </xs:sequence>
  </xs:complexType>
</xs:element>
```

```

</xs:annotation>
</xs:element>
<xs:element name="monthly" type="xs:int">
  <xs:annotation>
    <xs:documentation>
      <osd:label>Monthly payment </osd:label>
    </xs:documentation>
  </xs:annotation>
</xs:element>
<xs:element name="cost" type="xs:int">
  <xs:annotation>
    <xs:documentation>
      <osd:label>Cost</osd:label>
    </xs:documentation>
  </xs:annotation>
</xs:element>
</xs:sequence>
</xs:complexType>
</xs:element>

```

Aggregated lists have a dedicated editor in EBX®. This editor allows you to add or to delete occurrences.

### Attention

The addition of an `osd:table` declaration to an element with `maxOccurs > 1` is a very important consideration that must be taken into account during the design process. An aggregated list is severely limited with respect to the many features that are supported by tables. Some features unsupported on aggregated lists that are supported on tables are:

- Performance and memory optimization;
- Lookups, filters and searches;
- Sorting, view and display in hierarchies;
- Identity constraints (primary keys and uniqueness constraints);
- Detailed permissions for creation, modification, deletion and particular permissions at the record level;
- Detailed comparison and merge.

Thus, *aggregated lists should be used only for small volumes of simple data (one or two dozen occurrences), with no advanced requirements*. For larger volumes of data or more advanced functionalities, it is strongly advised to use an `osd:table` declaration.

For more information on table declarations, see [Tables and relationships](#) [p 495].

## 81.7 Including external data models

Including another data model in your current model allows you to use the reusable types that are defined in that data model. You can thus use the inclusion of external data models to share data types between multiple XML Schema Documents.

To include another XML Schema Document in your model, thereby including the data types that it defines, specify the `xs:include` element as follows:

```

<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:osd="urn:ebx-schemas:common_1.0" xmlns:fmt="urn:ebx-schemas:format_1.0">
  <xs:include schemaLocation="./schemaToInclude.xsd"/>
  ...
</xs:schema>

```

The attribute `schemaLocation` is mandatory and must specify either an absolute or a relative path to the XML Schema Document to include.

The inclusion of XML Schema Documents is not namespace aware, thus all included data types must belong to the same namespace. As a consequence, including XML Schema Documents that define data types of the same name is not supported.

EBX® includes extensions with specific URNs for including embedded data models and data models packaged in modules.

To include an embedded data model in a model, specify the URN defined by EBX®. For example:

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:osd="urn:ebx-schemas:common_1.0" xmlns:fmt="urn:ebx-schemas:format_1.0">
  <xs:include schemaLocation="urn:ebx:publication:myPublication"/>
  ...
</xs:schema>
```

To include a data model packaged in a module, specify the specific URN defined by EBX®. For example:

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:osd="urn:ebx-schemas:common_1.0" xmlns:fmt="urn:ebx-schemas:format_1.0">
  <xs:include schemaLocation="urn:ebx:module:aModuleName:/WEB-INF/ebx/schema/myDataModel.xsd"/>
  ...
</xs:schema>
```

See `SchemaLocation`<sup>API</sup> for more information about specific URNs supported by EBX®.

#### Note

If the packaged data model uses Java resources, the class loader of the module containing the data model will be used at runtime for resolving these resources.

## CHAPTER 82

# Tables and relationships

This chapter contains the following topics:

1. [Tables](#)
2. [Foreign keys](#)
3. [Associations](#)

## 82.1 Tables

### Overview

TIBCO EBX® supports the features of relational database tables, including the handling of large volumes of records, and identification by primary key.

Tables provide many benefits that are not offered by [aggregated lists](#) [p 492]. Beyond relational capabilities, some features that tables provide are:

- filters and searches;
- sorting, views and hierarchies;
- identity constraints: primary keys, [foreign keys](#) [p 500] and [uniqueness constraints](#) [p 517];
- specific permissions for creation, modification, and deletion;
- dynamic and contextual permissions at the individual record level;
- detailed comparison and merge;
- ability to have inheritance at the record level (see [dataset inheritance](#) [p 270]);
- performance and memory optimization.

#### See also

[Foreign keys](#) [p 500]

[Associations](#) [p 503]

[Working with existing datasets](#) [p 121]

[Simple tabular views](#) [p 114]

[Hierarchical views](#) [p 115]

[History](#) [p 249]

## ***Declaration***

A table element, which is an element with *maxOccurs* > 1, is declared by adding the following annotation:

```
<xs:annotation>
  <xs:appinfo>
    <osd:table>
      <primaryKey>/pathToField1 /pathToField...n</primaryKey>
    </osd:table>
  </xs:appinfo>
</xs:annotation>
```



## Common properties

| Element      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Required |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| primaryKeys  | <p>Specifies the primary key fields of the table.</p> <p>Each field of the primary key must be denoted by its absolute XPath notation that starts just under the root element of the table. If there are multiple fields in the primary key, the list is delimited by whitespace.</p> <p><b>Note:</b> Whitespaces in primary keys of type <code>xs:string</code> are handled differently. See <a href="#">Whitespace handling for primary keys of type string</a> [p 528].</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Yes.     |
| defaultLabel | <p>Defines the end-user display of records. Multiple variants can be specified:</p> <ul style="list-style-type: none"> <li>A static non-localized expression is defined using the <code>defaultLabel</code> element, for example: <pre>&lt;defaultLabel&gt;Product: \${./productCode}&lt;/defaultLabel&gt;</pre> </li> <li>Static localized expressions are specified using the <code>defaultLabel</code> element with the attribute <code>xml:lang</code>, for example: <pre>&lt;defaultLabel xml:lang="fr-FR"&gt;Produit : \${./productCode}&lt;/defaultLabel&gt; &lt;defaultLabel xml:lang="en-US"&gt;Product: \${./productCode}&lt;/defaultLabel&gt;</pre> </li> <li>A JavaBean that implements the interface <code>UILabelRenderer<sup>API</sup></code> and/or the interface <code>UILabelRendererForHierarchy<sup>API</sup></code>. The JavaBean is specified by means of the attribute <code>osd:class</code>, for example: <pre>&lt;defaultLabel osd:class="com.wombat.MyLabel"&gt;&lt;/defaultLabel&gt;</pre> </li> </ul> <p><b>Note:</b> The priority of the tags when displaying the user interface is the following:</p> <ol style="list-style-type: none"> <li><code>defaultLabel</code> tags with a JavaBean (but it is not allowed to define several renderers of the same type);</li> <li><code>defaultLabel</code> tags with a static localized expression using the <code>xml:lang</code> attribute;</li> <li><code>defaultLabel</code> tags with a static non-localized expression.</li> </ol> <p><b>Attention:</b> Access rights defined on associated datasets are not applied when displaying record labels. Fields that are usually hidden due to access rights restrictions will be displayed in labels.</p> | No.      |
| index        | <p>Specifies an index for speeding up requests that match this index (see <a href="#">performances</a> [p 293]).</p> <p>The attribute name is mandatory. Each field of the index must be denoted by its absolute XPath notation, which starts just under the root element of the table. If there are multiple fields in the index, the list is delimited by whitespace.</p> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>Indexing only concerns semantic and relational tables. History and replica tables are not affected.</li> <li>It is possible to define multiple indexes on a table.</li> <li>It is not possible to define two indexes with the same name.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | No.      |

| Element    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Required |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
|            | <ul style="list-style-type: none"> <li>It is not possible to declare two indexes containing the exact same fields.</li> <li>An indexed field must be terminal.</li> <li>An indexed field cannot be a list nor under a list.</li> <li>A field declared as an <i>inherited field</i> cannot be indexed.</li> <li>A field declared as a function cannot be indexed.</li> </ul> <p>For performance purposes, the following nodes are automatically indexed:</p> <ul style="list-style-type: none"> <li>Primary keys nodes. See <a href="#">primary keys</a> [p 495].</li> <li>Nodes defining a foreign key constraint. See <a href="#">foreign key constraint</a> [p 500].</li> <li>Nodes declared as being unique. See <a href="#">uniqueness constraint</a> [p 517].</li> <li>Auto-incremented nodes. See <a href="#">auto-incremented values</a> [p 531].</li> </ul> |          |
| recordForm | <p>Defines a specific component for customizing the record form in a dataset. This component is defined using a JavaBean that extends <code>UIForm<sup>api</sup></code> or implements <code>UserServiceRecordFormFactory<sup>api</sup></code>.</p> <p>The JavaBean is specified by means of the attribute <code>osd:class</code>, for example:</p> <pre>&lt;recordForm osd:class="com.wombat.MyRecordForm"/&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                               | No.      |

## Example

Below is an example of a product catalog:

```
<xs:element name="Products" minOccurs="0" maxOccurs="unbounded">
  <xs:annotation>
    <xs:documentation>
      <osd:label>Product Table </osd:label>
      <osd:description>List of products in Catalog </osd:description>
    </xs:documentation>
  </xs:annotation>
  <xs:complexType>
    <xs:annotation>
      <xs:appinfo>
        <osd:table>
          <primaryKeys>./productRange /productCode</primaryKeys>
          <index name="indexProductCode">/productCode</index>
        </osd:table>
      </xs:appinfo>
    </xs:annotation>
    <xs:sequence>
      <xs:element name="productRange" type="xs:string"/><!-- key -->
      <xs:element name="productCode" type="xs:string"/><!-- key -->
      <xs:element name="productLabel" type="xs:string"/>
      <xs:element name="productDescription" type="xs:string"/>
      <xs:element name="productWeight" type="xs:int"/>
      <xs:element name="productType" type="xs:string"/>
      <xs:element name="productCreationDate" type="xs:date"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>

<xs:element name="Catalogs" minOccurs="0" maxOccurs="unbounded">
  <xs:annotation>
    <xs:documentation>
      <osd:label>Catalog Table</osd:label>
      <osd:description>List of catalogs</osd:description>
    </xs:documentation>
  </xs:annotation>
  <xs:complexType>
    <xs:annotation>
      <xs:appinfo>
        <osd:table>
          <primaryKeys>/catalogId</primaryKeys>
        </osd:table>
      </xs:appinfo>
    </xs:annotation>
  </xs:complexType>
</xs:element>
```

```

</osd:table>
</xs:appinfo>
</xs:annotation>
<xs:sequence>
  <xs:element name="catalogId" type="xs:string"/><!-- key -->
  <xs:element name="catalogLabel" type="xs:string"/>
  <xs:element name="catalogDescription" type="xs:string"/>
  <xs:element name="catalogType" type="xs:string"/>
  <xs:element name="catalogPublicationDate" type="xs:date"/>
</xs:sequence>
</xs:complexType>
</xs:element>

```

## Properties related to dataset inheritance

The following properties are only valid in the context of dataset inheritance:

| Element                          | Description                                                                                                                                                           | Required               |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| onDelete-deleteOccultingChildren | Specifies whether, upon record deletion, child records in occulting mode are also to be deleted.<br>Valid values are: never or always.                                | No, default is never.  |
| mayCreateRoot                    | Specifies whether root record creation is allowed. The expression must follow the syntax below. See <a href="#">definition modes</a> [p 270].                         | No, default is always. |
| mayCreateOverwriting             | Specifies whether records are allowed to be overwritten in child datasets. The expression must follow the syntax below. See <a href="#">definition modes</a> [p 270]. | No, default is always. |
| mayCreateOcculting               | Specifies whether records are allowed to be occulted in child datasets. The expression must follow the syntax below. See <a href="#">definition modes</a> [p 270].    | No, default is always. |
| mayDuplicate                     | Specifies whether record duplication is allowed. The expression must follow the syntax below.                                                                         | No, default is always. |
| mayDelete                        | Specifies whether record deletion is allowed. The expression must follow the syntax below.                                                                            | No, default is always. |

The may... expressions specify when the action is possible, though the ultimate availability of the action also depends on the user access rights. The expressions have the following syntax:

expression ::= always | never | <condition>\*

condition ::= [root:yes | root:no]

"always": the operation is "always" possible (but user rights may restrict this).

"never": the operation is never possible.

"root:yes": the operation is possible if the record is in a root instance.

"root:no": the operation is not possible if the record is in a root instance.

If the record does not define any specific conditions, the default is used.

See also [Dataset inheritance](#) [p 269]

## Using toolbars

It is possible to define the toolbars to display in the user interface using the element `defaultView/toolbars` under `xs:annotation/appinfo/osd:table`. A toolbar allows to customize the buttons and menus to display when displaying a table view, a hierarchical view, or a record form.

The table below presents the elements that can be defined under `defaultView/toolbars`.

| Element                       | Description                                                            | Required |
|-------------------------------|------------------------------------------------------------------------|----------|
| <code>tabularViewTop</code>   | Defines the toolbar to use on top of the default table view.           | No.      |
| <code>tabularViewRow</code>   | Defines the toolbar to use on each row of the default table view.      | No.      |
| <code>recordTop</code>        | Defines the toolbar to use in the record form.                         | No.      |
| <code>hierarchyViewTop</code> | Defines the toolbar to use in the default hierarchy view of the table. | No.      |

See also [Toolbars](#) [p 549]

## Example

Below is an example of custom toolbars used by a product catalog:

```
<xs:element name="Products" minOccurs="0" maxOccurs="unbounded">
  <xs:annotation>
    <xs:documentation>
      <osd:label>Product Table </osd:label>
      <osd:description>List of products in Catalog </osd:description>
    </xs:documentation>
  </xs:annotation>
  <xs:complexType>
    <xs:annotation>
      <xs:appinfo>
        <osd:table>
          <primaryKeys>./productRange /productCode</primaryKeys>
          <defaultView>
            <toolbars>
              <tabularViewTop>toolbar_name_for_tabularViewTop</tabularViewTop>
              <tabularViewRow>toolbar_name_for_tabularViewRow</tabularViewRow>
              <recordTop>toolbar_name_for_recordTop</recordTop>
              <hierarchyViewTop>toolbar_name_for_hierarchyViewTop</hierarchyViewTop>
            </toolbars>
          </defaultView>
        </osd:table>
      </xs:appinfo>
    </xs:annotation>
    ...
  </xs:complexType>
</xs:element>
```

### Note

If a toolbar does not exist or is not available for a specific location then no toolbar will be displayed in the user interface in the corresponding location.

## 82.2 Foreign keys

### Declaration

A reference to a [table](#) [p 495] is defined using the extended facet `osd:tableRef`.

The node holding the `osd:tableRef` declaration must be of type `xs:string`. At the instantiation, any value of the node identifies a record in the target table using its **primary key syntax** `PrimaryKey`.

syntax<sup>API</sup>. This extended facet is also interpreted as an enumeration whose values refer to the records in the target table.

| Element   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Required                                                                                                                 |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| tablePath | XPath expression that specifies the target table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Yes.                                                                                                                     |
| container | Reference of the dataset that contains the target table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Only if the dataspace element is defined. Otherwise, default is the current dataset.                                     |
| branch    | Reference of the dataspace that contains the container dataset.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | No, default is the current dataspace or snapshot.                                                                        |
| display   | <p>Custom display for presenting the selected foreign key in the current record and the sorted list of possible keys. Two variants can be specified, either pattern-based expressions, or a JavaBean if the needs are very specific:</p> <ul style="list-style-type: none"> <li>Static expressions are specified using the <code>display</code> and <code>pattern</code> elements. These static expressions can be localized using the additional attribute <code>xml:lang</code> on the <code>pattern</code> element, for example: <pre>&lt;display&gt; &lt;pattern&gt;Product : \${../productCode}&lt;/pattern&gt; &lt;pattern xml:lang="fr-FR"&gt;Produit : \${../productCode}&lt;/pattern&gt; &lt;pattern xml:lang="en-US"&gt;Product: \${../productCode}&lt;/pattern&gt; &lt;/display&gt;</pre> </li> <li>A JavaBean that implements the interface <code>TableRefDisplay</code><sup>API</sup>. It is specified using the attribute <code>osd:class</code>. For example: <pre>&lt;display osd:class="com.wombat.MyLabel"&gt;&lt;/display&gt;</pre> </li> </ul> <p>It is not possible to define both variants on the same foreign key element.</p> <p><b>Attention:</b> Access rights defined on associated datasets are not applied when displaying record labels. Fields that are usually hidden due to access rights restrictions will be displayed in labels.</p> | No, if the <code>display</code> property is not specified, the <a href="#">table's record rendering</a> (p 497) is used. |
| filter    | <p>Specifies an additional constraint that filters the records of the target table. Two types of filters are available:</p> <ul style="list-style-type: none"> <li>An XPath filter is an XPath predicate in the target table context. It is specified using the <code>predicate</code> element. For example: <pre>&lt;filter&gt;&lt;predicate&gt;type = \${../refType}&lt;/predicate&gt;&lt;/filter&gt;</pre> </li> </ul> <p>A localized validation message can be specified using the element <code>validationMessage</code>, which will be displayed to the end-user at the validation time if a record is not accepted by the filter.</p> <p>A specific severity level can be defined in a nested <code>severity</code> element. The default severity is 'error'.</p> <p>Each localized message variant is defined in a nested <code>message</code> element with its locale in an <code>xml:lang</code> attribute.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                | No.                                                                                                                      |

| Element    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Required |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
|            | <p>To specify a default message for unsupported locales, define a message element with no <code>xml:lang</code> attribute.</p> <p>In the user interface, the XPath filter is applied to filter a table according to the value of a foreign key field. That is, if a foreign key field specifies an XPath filter in a data model, then it will be reused in the filter pane to restrict the set of values in the associated combo-box displayed in the filter pane. However, the predicate used by the filter pane will only take into account the non-contextual parts of the predicate.</p> <ul style="list-style-type: none"> <li>A programmatic filter is a JavaBean that implements the interface <code>TableRefFilter<sup>API</sup></code>. It is specified using the attribute <code>osd:class</code>. For example:</li> </ul> <pre>&lt;filter osd:class="com.wombat.MyFilter"&gt;&lt;/filter&gt;</pre> <p>Additional validation messages can be specified during the setup of the programmatic filter.</p> <p>In the user interface, programmatic filters are not applied to filter the set of values in the associated combo-box displayed in the filter pane. However, it is possible to specify an additional XPath predicate that will be used in the filter pane of the user interface. This XPath predicate is specified during the setup of the programmatic filter using the method <code>TableRefFilterContext.setFilterForSearch<sup>API</sup></code>.</p> <p><b>Note:</b></p> <p>The attributes <code>osd:class</code> and the property <code>predicate</code> cannot be set simultaneously.</p> <p>The validation search XPath functions are forbidden on a <code>tableRef</code> filter.</p> |          |
| validation | <p>Specifies localized validation messages for the <code>osd:tableRef</code> and error management policy.</p> <p>A specific severity level can be defined in a nested <code>severity</code> element. The default severity is 'error'.</p> <p>An error management policy can be defined in a nested <code>blocksCommit</code> element. The error management policy that blocks all operations does not apply to filters. That is, a foreign key constraint is not blocking if a referenced record exists but does not satisfy a foreign key filter. In this case, updates are not rejected, and a validation error will be reported.</p> <p>Each localized message variant is defined in a nested <code>message</code> element with its locale in an <code>xml:lang</code> attribute. To specify a default message for unsupported locales, define a message element with no <code>xml:lang</code> attribute.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | No.      |

### Attention

You can create a dataset which has a foreign key to a container that does not exist in the repository. However, the content of this dataset will not be available until the container is created. After the creation of the container, a data model refresh is required to make the dataset available. When creating a dataset that refers to a container that does not yet exist, the following limitations apply:

- Triggers defined at the dataset level are not executed.
- Default values for fields that are not contained in tables are not initialized.

- During an archive import, it is not possible to create a dataset that refers to a container that does not exist.

## Example

The example below specifies a foreign key in the 'Products' table to a record of the 'Catalogs' table.

```
<xs:element name="catalog_ref" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:tableRef>
          <tablePath>/root/Catalogs</tablePath>
          <display>
            <pattern xml:lang="en-US">Catalog: ${./catalogId}</pattern>
            <pattern xml:lang="fr-FR">Catalogue : ${./catalogId}</pattern>
          </display>
          <validation>
            <severity>error</severity>
            <blocksCommit>onInsertUpdateOrDelete</blocksCommit>
            <message>A default error message</message>
            <message xml:lang="en-US">A localized error message</message>
            <message xml:lang="fr-FR">Un message d'erreur localisé</message>
          </validation>
        </osd:tableRef>
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

### See also

[Table definition](#) [p 495]

**Primary key syntax** *PrimaryKey.syntax*<sup>API</sup>

[Extraction of foreign keys \(XPath predicate syntax\)](#) [p 230]

[Associations](#) [p 503]

[View for advanced selection](#) [p 543]

*SchemaNode.getFacetOnTableReference*<sup>API</sup>

*SchemaFacetTableRef*<sup>API</sup>

## 82.3 Associations

### Overview

An association provides an abstraction over an existing relationship in the data model, and allows an easy model-driven integration of *associated objects* in the user interface and in data services.

Several types of associations are supported:

- 'By foreign key' specifies the inverse relationship of an existing [foreign key field](#) [p 500].
- 'Over a link table' specifies a relationship based on an intermediate link table (such tables are often called "join tables"). This link table has to define two foreign keys, one referring to the 'source' table (the table holding the association element) and another one referring to the 'target' table.
- 'By an XPath predicate' specifies a relationship based on an XPath predicate.

For an association, it is also possible to:

- Filter associated objects by specifying an additional XPath filter.

- Configure a tabular view to define the fields that must be displayed in the associated table.
- Define how associated objects are to be rendered in forms.
- Hide/show associated objects in the data service 'select' operation. See [Hiding a field in Data Services](#) [p 542].
- Specify the minimum and maximum number of associated objects that are required.
- Add validation constraints using XPath predicates for restricting associated objects.

#### See also

*SchemaNode.getAssociationLink*<sup>API</sup>

*SchemaNode.isAssociationNode*<sup>API</sup>

*AssociationLink*<sup>API</sup>

## Declaration

Associations are defined in the data model using the XML Schema element `osd:association` under `xs:annotation/appInfo`.

#### Restrictions:

- An association must be a simple element of type *xs:string*.
- An association can only be defined inside a table.

#### Note

The "official" cardinality constraints (`minOccurs="0"` `maxOccurs="0"`) are required because, from an instance of XML Schema, the corresponding node is absent. In other words, an association has no value and is considered as a "virtual" element as far as XML and XML Schema is concerned.



The table below presents the elements that can be defined under `xs:annotation/appInfo/osd:association`.

| Element         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Required                                                                                   |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| tableRefInverse | <p>Defines the properties of an association that is the inverse relationship of a <i>foreign key</i>.</p> <p>Element <code>fieldToSource</code> defines the foreign key that refers to the source table of the association. The element <code>fieldToSource</code> is mandatory and must specify a foreign key field that refers to the table containing the association.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Yes if the association is the inverse relationship of a <i>foreign key</i> , otherwise no. |
| linkTable       | <p>Defines the properties of an association over a link table.</p> <p>The element <code>table</code> specifies the link table used by the association. The element <code>table</code> is mandatory and must refer to an existing table.</p> <p><b>Important:</b> In order to be used by an association, a link table must define a primary key that is composed of auto-incremented fields and/or the foreign key to the source or target table of the association.</p> <p>The element <code>fieldToSource</code> defines the foreign key that refers to the source table of the association. The element <code>fieldToSource</code> is mandatory and must specify a foreign key field that refers to the table containing the association.</p> <p>The element <code>fieldToTarget</code> defines the foreign key that refers to the target table of the association. The element <code>fieldToTarget</code> is mandatory and must specify a foreign key field.</p>                                                                                                                                             | Yes if the association is over a link table, otherwise no.                                 |
| xpathLink       | <p>Defines the properties of an association that is based on an <i>XPath predicate</i>.</p> <p>The predicate element specifies the criteria of the association, relative to the current node.</p> <p>Examples: <code>/root/Products[catalog_ref = \${../catalogId}]</code> or <code>//Products[catalog_ref = \${../catalogId}]</code> or <code>../Products[catalog_ref = \${../catalogId}]</code>.</p> <p>The path to the predicate, for example <code>../Products</code>, specifies the target table of the association. This part of the path is resolved with respect to the current table. It is not possible to refer to a table using a relative path if the association targets a table in another dataset.</p> <p>If the association depends on fields of the source table, the XPath expression predicate must include references to the elements on which it depends using the notation <code>\${&lt;relative-path&gt;}</code> where <i>relative-path</i> is a path that identifies the element relative to the association node.</p> <p>See <a href="#">EBX® XPath supported syntax</a> [p 225].</p> | Yes if the association is based on an <i>XPath predicate</i> , otherwise no.               |

Note

| Element     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Required |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
|             | <p>The validation search XPath functions are forbidden on an XPath link.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |          |
| filter      | <p>Defines an XPath predicate to filter associated objects using the predicate element. For example:</p> <pre>&lt;filter&gt;&lt;predicate&gt;type = \${../refType}&lt;/predicate&gt;&lt;/filter&gt;</pre> <p>It is only possible to use fields from the source and the target tables when defining an XPath filter. That is, if it is an association over a link table, it is not possible to use fields of the link table in the XPath filter.</p> <p><b>Error message on creation:</b> in the user interface, the record creation is blocked when a user submits a new associated record that does not comply with the filter. The error message can be customized using the element <code>checkOnAssociatedRecordCreation/message</code>. Each localized message variant is defined in a nested message element with its locale in an <code>xml:lang</code> attribute. To specify a default message for unsupported locales, define a message element with no <code>xml:lang</code> attribute. See <a href="#">Examples</a> (p 512) for more information on this property.</p> <p><b>Note</b></p> <p>The validation search XPath functions are forbidden for association filter.</p> | No.      |
| xpathFilter | <p><b>Note:</b> Deprecated. This property has been replaced by the property <code>filter</code>.</p> <p>Defines an XPath predicate to filter associated objects.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | No.      |
| recordForm  | <p>Defines a specific component for customizing the form of an associated record. This component is defined using a JavaBean that implements <code>UserServiceAssociationRecordFormFactory<sup>API</sup></code>.</p> <p>The JavaBean is specified by means of the attribute <code>osd:class</code>, for example:</p> <pre>&lt;recordForm osd:class="com.wombat.MyRecordFormFactory"/ &gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | No.      |

It is possible to refer to another dataset. For that, the following properties must be defined either under the element `tableRefInverse`, `linkTable` or `xpathLink` depending on the type of the association:

| Element                     | Description                                                                                                                                                                                                                                                                                                       | Required |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| <code>schemaLocation</code> | Defines the data model containing the fields used by the association. The data model is defined using a specific URN that allows referring to embedded data models and data models packaged in modules. See <code>SchemaLocation<sup>API</sup></code> for more information about specific URNs supported by EBX®. | Yes.     |
| <code>dataSet</code>        | Defines the dataset used by the association. This dataset must use the data model specified by the element <code>schemaLocation</code> .                                                                                                                                                                          | Yes.     |
| <code>dataSpace</code>      | Defines the dataspace containing the dataset used by the association.                                                                                                                                                                                                                                             | No.      |

**Important:** When creating a dataset, you can create a dataset that defines an association to a container that does not yet exist in the repository. However, the content of this dataset will not be available immediately upon creation. After the absent container is created, a data model refresh is required in order to make the dataset available. When creating a dataset that refers to a container that does not yet exist, the following limitations apply:

- Triggers defined at the dataset level are not executed.
- Default values on fields outside tables are not initialized.
- During an archive import, it is not possible to create a dataset that refers to a container that does not exist.

## User interface integration

It is possible to define how associated objects are to be rendered in forms, using the element `osd:defaultView/displayMode` under `xs:annotation/appinfo`.

Possible values are:

- `inline`, specifies that associated records are to be rendered in the form at the same position of the association in the data model.
- `tab`, specifies that associated records are to be rendered in a specific tab.
- `link`, specifies that associated records are to be rendered in a modal window.

By default, associated records are rendered `inline` if this property is not defined.

The following example specifies that associated objects are to be rendered `inline` in the form:

```
<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
      </osd:association>
      <osd:defaultView>
        <displayMode>inline</displayMode>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
```

</xs:element>

The following example specifies that associated objects are to be rendered in a specific tab:

```
<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
      </osd:association>
      <osd:defaultView>
        <displayMode>tab</displayMode>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

### Using toolbars

It is possible to define the toolbars to display in the user interface using the element `osd:defaultView/toolbars` under `xs:annotation/appinfo`. A toolbar allows to customize the buttons and menus to display when displaying the tabular view of an association.

The table below presents the elements that can be defined under `osd:defaultView/toolbars`.

| Element        | Description                                                                      | Required |
|----------------|----------------------------------------------------------------------------------|----------|
| tabularViewTop | Defines the toolbar to use in the default table view of this association.        | No.      |
| tabularViewRow | Defines the toolbar to use for each row of the default view of this association. | No.      |

The following example shows how to use toolbars from the previous association between a catalog and its products:

```
<xs:element name="Products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
      </osd:association>
      <osd:defaultView>
        <toolbars>
          <tabularViewTop>toolbar_name_for_tabularViewTop</tabularViewTop>
          <tabularViewRow>toolbar_name_for_tabularViewRow</tabularViewRow>
        </toolbars>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

#### Note

It is only possible to use the toolbars defined in the data model containing the target table of the association. That is, if the target table of the association is defined in another data model, then it is only possible to reference a toolbar defined in this data model and not in the one holding the association.

See also [Toolbars](#) [p 549]

## Customized view of associated objects

A specific tabular view can be specified to define the fields that must be displayed in the target table. If a tabular view is not defined, all columns that a user is allowed to view, according to the granted access rights, are displayed. A tabular view is defined using the element `osd:defaultView/tabularView` under `xs:annotation/appinfo`.

The table below shows the elements that can be defined under `osd:defaultView/tabularView`.

| Element | Description                                                                                                                                                                                                                                                                                                                                                                                                     | Required |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| column  | Define a field of the target table to display. The specified path must be absolute from the target table and must refer to an existing field. Several column elements can be defined to specify the fields that are to be displayed.                                                                                                                                                                            | No.      |
| sort    | Define a field that can be used to sort associated objects. Several sort elements can be defined to specify the fields that can be used to sort associated objects.<br><br>The element <code>nodePath</code> defines the path of the field that can be used to sort associated objects.<br><br>The element <code>isAscending</code> specifies whether the sort order is ascending (true) or descending (false). | No.      |

The following example shows how to define a tabular view from the previous association between a catalog and its products:

```
<xs:element name="Products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
      </osd:association>
      <osd:defaultView>
        <tabularView>
          <column>/productRange</column>
          <column>/productCode</column>
          <column>/productLabel</column>
          <column>/productDescription</column>
          <sort>
            <nodePath>/productLabel</nodePath>
            <isAscending>true</isAscending>
          </sort>
        </tabularView>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

## Actions in the user interface

In the user interface, it is possible to perform the following actions:

- **Create:** it allows directly creating an object in the target table of the association. When a new object is created, it is automatically associated with the current record.
- **Duplicate:** allows to duplicate an object in the target table of the association. When a new object is created, it is automatically associated with the current record.

- **Associate:** associates an existing object with the current record. In the case of an association over a link table, a record in the link table is automatically created to materialize the link between the current record and the existing object.
- **Move:** associates the selected objects to a different record than the current one. In the case of an association over a link table, the previous link record is automatically deleted and a new record in the link table is automatically created to materialize the link between the selected objects and their new parent record.
- **Delete:** deletes selected associated objects in the target table of the association.
- **Detach:** breaks the semantic link between the current record and the selected associated objects. In the case of an association over a link table, the records in the link table are automatically deleted, to break the links between the current record and associated objects.

#### Note

The actions *associate* and *detach* are not available when the association is defined using an **XPath predicate** (element *xpathLink*).

### Customized view for actions

A published view, tabular or hierarchical, can be specified to define how objects should be displayed when performing an action through the user interface. A published view is defined using the element `osd:defaultView/associationViews` under `xs:annotation/appinfo`.

The table below shows the elements that can be defined under `osd:defaultView/associationViews`.

| Element                             | Description                                                                                                                                                                                                             | Required |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| <code>viewForAssociateAction</code> | Define a published view to be used when displaying the objects in the target table to be associated with the current record. The specified view must be published and created upon the target table of the association. | No.      |
| <code>viewForMoveAction</code>      | Define a published view to be used when moving an associated object to another record of the current table. The specified view must be published and created upon the current table.                                    | No.      |

The following example shows how to define views from the previous association between a catalog and its products:

```
<xs:element name="Products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
      </osd:association>
      <osd:defaultView>
        <associationViews>
          <viewForAssociateAction>view_name_for_catalogs</viewForAssociateAction>
          <viewForMoveAction>view_name_for_products</viewForMoveAction>
        </associationViews>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

## Validation

Some controls can be defined on associations, in order to restrict associated objects. These controls are defined under the element `osd:association`.

The table below presents the controls that can be defined under `xs:annotation/appInfo/osd:association`.

| Element                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Required                                      |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| <code>minOccurs</code>  | Specifies the minimum number of associated objects that are required for this association. This minimum number is defined using the element <code>value</code> and must be a positive integer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | No, by default the minimum is not restricted. |
| <code>maxOccurs</code>  | Specifies the maximum number of associated objects that are allowed for this association. This maximum number is defined by the element <code>value</code> and must be either a positive integer or the raw string <code>unbounded</code> which indicates that this maximum is not restricted. The maximum number of associated objects must be greater than the minimum number of associated objects.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | No, by default the maximum is not restricted. |
| <code>constraint</code> | <p>Defines an XPath predicate for restricting associated records. It is only possible to use fields from the source and the target table when defining an XPath predicate. That is, if it is an association over a link table, it is not possible to use fields of the link table in the XPath predicate.</p> <p>In associated datasets, a validation message of the specified severity is added and displayed to the end-user at the validation time when an associated record does not comply with the specified constraint.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       | No.                                           |
| <code>validation</code> | <p>A validation message can be defined under the elements <code>minOccurs</code>, <code>maxOccurs</code> and <code>constraint</code>, using the element <code>validation</code>. The severity of the validation message is specified using the element <code>severity</code>. Possible severities are: <code>error</code>, <code>warning</code> and <code>info</code>.</p> <p>If the severity is not specified then, by default, the severity <code>error</code> is used.</p> <p>A localized validation message can be specified using the element <code>message</code>, which will be displayed to the end-user at the validation time if an association does not comply with this constraint. Each localized message variant is defined in a nested message element with its locale in an <code>xml:lang</code> attribute. To specify a default message for unsupported locales, define a message element with no <code>xml:lang</code> attribute.</p> | No.                                           |

## Data services integration

It is possible to define whether associated objects must be hidden in the Data service `select` operation. For this, the property `osd:defaultview/hiddenInDataServices` under `xs:annotation/xs:appinfo` can be set on the association. Setting the property to 'true' will hide associated objects in the Data

service select operation. If this property is not defined then, by default, associated objects will be shown in the Data service select operation.

### See also

[Hiding a field in Data Services](#) [p 542]

[Association field](#) [p 618]

## Examples

For example, the product catalog data model defined [previously](#) [p 498] specifies that a product belongs to a catalog (explicitly defined by a foreign key in the 'Products' table). The reverse relationship (that a catalog has certain products) is not easily represented in XML Schema, unless the 'Catalogs' table includes the following association that is the inverse of a foreign key:

```
<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
      </osd:association>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

For an association over a link table, we can consider the previous example and bring some updates. For instance, the foreign key in the 'Products' table is deleted and the relation between a product and a catalog is redefined by a link table (named 'Catalogs\_Products') that has a primary key composed of two foreign keys: one that refers to the 'Products' table (named 'productRef') and another to the 'Catalogs' table (named 'catalogRef'). The following example shows how to define an association over a link table from this new relationship:

```
<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <linkTable>
          <table>/root/Catalogs_Products</table>
          <fieldToSource>./catalogRef</fieldToSource>
          <fieldToTarget>./productRef</fieldToTarget>
        </linkTable>
      </osd:association>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

The following example shows an association that refers to a foreign key in another dataset. In this example, the 'Products' and 'Catalogs' tables are not in the same dataset:

```
<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <schemaLocation>urn:ebx:module:aModuleName:/WEB-INF/ebx/schema/products.xsd</schemaLocation>
          <dataSet>Products</dataSet>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
      </osd:association>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

The following example defines an XPath filter to associate only products of the 'Technology' type:

```
<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
```



```

<osd:association>
<tableRefInverse>
  <fieldToSource>/root/Products/catalog_ref</fieldToSource>
</tableRefInverse>
<filter>
  <predicate>./productType = 'Technology'</predicate>
  <checkOnAssociatedRecordCreation>
    <message>A default message</message>
    <message xml:lang="en-US">A localized message</message>
    <message xml:lang="fr-FR">Un message localisé</message>
  </checkOnAssociatedRecordCreation>
</filter>
</osd:association>
</xs:appinfo>
</xs:annotation>
</xs:element>

```

The following example specifies the minimum number of products that are required for a catalog:

```

<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
        <minOccurs>
          <value>1</value>
        <validation>
          <severity>warning</severity>
          <message xml:lang="en-US">One product should at least be associated to this catalog.</message>
          <message xml:lang="fr-FR">Un produit doit au moins être associé à ce catalogue.</message>
        </validation>
        </minOccurs>
      </osd:association>
    </xs:appinfo>
  </xs:annotation>
</xs:element>

```

The following example specifies that a catalog must contain at most ten products:

```

<xs:element name="products" minOccurs="0" maxOccurs="0" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:association>
        <tableRefInverse>
          <fieldToSource>/root/Products/catalog_ref</fieldToSource>
        </tableRefInverse>
        <maxOccurs>
          <value>10</value>
        <validation>
          <severity>warning</severity>
          <message xml:lang="en-US">Too much products for this catalog.</message>
          <message xml:lang="fr-FR">Ce catalogue a trop de produits.</message>
        </validation>
        </maxOccurs>
      </osd:association>
    </xs:appinfo>
  </xs:annotation>
</xs:element>

```



## CHAPTER 83

# Constraints, triggers and functions

Facets allow you to define data constraints in your data models. TIBCO EBX® supports XML Schema facets and provides extended and programmatic facets for advanced data controls.

This chapter contains the following topics:

1. [XML Schema supported facets](#)
2. [Extended facets](#)
3. [Programmatic facets](#)
4. [Control policy](#)
5. [Triggers and functions](#)

## 83.1 XML Schema supported facets

The tables below show the facets that are supported by different data types.

**Key:**

- **X** - Supported
- **1** - The `whiteSpace` facet can be defined, but is not interpreted by EBX®
- **2** - In XML Schema, boundary facets are not allowed on the type `string`. Nevertheless, EBX® allows such facets as extensions.
- **3** - The `osd:resource` type only supports the facet `Facet0Resource`, which is required. See [Extended Facets](#) [p 518].

- **4** - `osd:dataspaceKey`, `osd:datasetName` and `osd:color` types do not support facets. Only [Programmatic constraints](#) [p 522] are supported on these types.

|                                                       | <a href="#">length</a> | <a href="#">minLength</a> | <a href="#">max Length</a> | <a href="#">pattern</a> | <a href="#">enumeration</a> | <a href="#">white Space</a> |
|-------------------------------------------------------|------------------------|---------------------------|----------------------------|-------------------------|-----------------------------|-----------------------------|
| <a href="#">xs:string</a>                             | X                      | X                         | X                          | X                       | X                           | 1                           |
| <a href="#">xs:boolean</a>                            |                        |                           |                            | X                       |                             | 1                           |
| <a href="#">xs:decimal</a>                            |                        |                           |                            | X                       | X                           | 1                           |
| <a href="#">xs:dateTime</a>                           |                        |                           |                            | X                       | X                           | 1                           |
| <a href="#">xs:time</a>                               |                        |                           |                            | X                       | X                           | 1                           |
| <a href="#">xs:date</a>                               |                        |                           |                            | X                       | X                           | 1                           |
| <a href="#">xs:anyURI</a>                             | X                      | X                         | X                          | X                       | X                           | 1                           |
| <a href="#">xs:Name</a>                               | X                      | X                         | X                          | X                       | X                           | 1                           |
| <a href="#">xs:integer</a>                            |                        |                           |                            | X                       | X                           | 1                           |
| <a href="#">osd:resource</a> [p 484] <sup>3</sup>     |                        |                           |                            |                         |                             | 1                           |
| <a href="#">osd:dataspaceKey</a> [p 486] <sup>4</sup> |                        |                           |                            |                         |                             | 1                           |
| <a href="#">osd:datasetName</a> [p 488] <sup>4</sup>  |                        |                           |                            |                         |                             | 1                           |
| <a href="#">osd:color</a> [p 486] <sup>4</sup>        |                        |                           |                            |                         |                             | 1                           |

|                             | <a href="#">fraction Digits</a> | <a href="#">total Digits</a> | <a href="#">max Inclusive</a> | <a href="#">max Exclusive</a> | <a href="#">min Inclusive</a> | <a href="#">min Exclusive</a> |
|-----------------------------|---------------------------------|------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
| <a href="#">xs:string</a>   |                                 |                              | 2                             | 2                             | 2                             | 2                             |
| <a href="#">xs:boolean</a>  |                                 |                              |                               |                               |                               |                               |
| <a href="#">xs:decimal</a>  | X                               | X                            | X                             | X                             | X                             | X                             |
| <a href="#">xs:dateTime</a> |                                 |                              | X                             | X                             | X                             | X                             |
| <a href="#">xs:time</a>     |                                 |                              | X                             | X                             | X                             | X                             |
| <a href="#">xs:date</a>     |                                 |                              | X                             | X                             | X                             | X                             |

|                                                       | <a href="#">fraction</a><br><a href="#">Digits</a> | <a href="#">total</a><br><a href="#">Digits</a> | <a href="#">max</a><br><a href="#">Inclusive</a> | <a href="#">max</a><br><a href="#">Exclusive</a> | <a href="#">min</a><br><a href="#">Inclusive</a> | <a href="#">min</a><br><a href="#">Exclusive</a> |
|-------------------------------------------------------|----------------------------------------------------|-------------------------------------------------|--------------------------------------------------|--------------------------------------------------|--------------------------------------------------|--------------------------------------------------|
| <a href="#">xs:anyURI</a>                             |                                                    |                                                 |                                                  |                                                  |                                                  |                                                  |
| <a href="#">xs:Name</a>                               |                                                    |                                                 | 2                                                | 2                                                | 2                                                | 2                                                |
| <a href="#">xs:integer</a>                            | X                                                  | X                                               | X                                                | X                                                | X                                                | X                                                |
| <a href="#">osd:resource</a> [p 484] <sup>3</sup>     |                                                    |                                                 |                                                  |                                                  |                                                  |                                                  |
| <a href="#">osd:dataspaceKey</a> [p 486] <sup>4</sup> |                                                    |                                                 |                                                  |                                                  |                                                  |                                                  |
| <a href="#">osd:datasetName</a> [p 488] <sup>4</sup>  |                                                    |                                                 |                                                  |                                                  |                                                  |                                                  |
| <a href="#">osd:color</a> [p 486] <sup>4</sup>        |                                                    |                                                 |                                                  |                                                  |                                                  |                                                  |

Example:

```
<xs:element name="loanRate">
  <xs:simpleType>
    <xs:restriction base="xs:decimal">
      <xs:minInclusive value="4.5" />
      <xs:maxExclusive value="17.5" />
    </xs:restriction>
  </xs:simpleType>
</xs:element>
```

## Uniqueness constraint

It is possible to define a uniqueness constraint, using the standard XML Schema element [xs:unique](#). This constraint indicates that a value or a set of values has to be unique inside a table.

**Example:**

In the example below, a uniqueness constraint is defined on the 'publisher' table, for the target field 'name'. This means that no two records in the 'publisher' table can have the same name.

```
<xs:element name="publisher">
  ...
  <xs:complexType>
    <xs:sequence>
      ...
      <xs:element name="name" type="xs:string" />
      ...
    </xs:sequence>
  </xs:complexType>
  <xs:unique name="uniqueName">
    <xs:annotation>
      <xs:appinfo>
        <osd:validation>
          <severity>error</severity>
          <message>Name must be unique in table.</message>
          <message xml:lang="en-US">Name must be unique in table.</message>
          <message xml:lang="fr-FR">Le nom doit être unique dans la table.</message>
        </osd:validation>
      </xs:appinfo>
    </xs:annotation>
    <xs:selector xpath="." />
    <xs:field xpath="name" />
  </xs:unique>
</xs:element>
```

A uniqueness constraint has to be defined within a table and has the following properties:

| Property            | Description                                                                                                                                                                                                                | Mandatory |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| name attribute      | Identifies the constraint in the data model.                                                                                                                                                                               | Yes       |
| xs:selector element | Indicates the table to which the uniqueness constraint applies using a restricted XPath expression ('..' is forbidden). It can also indicate an element within the table (without changing the meaning of the constraint). | Yes       |
| xs:field element    | Indicates the field in the context whose values must be unique, using a restricted XPath expression.<br>It is possible to indicate that a set of values must be unique by defining multiple xs:field elements.             | Yes       |

#### Note

Undefined values (null values) are ignored on uniqueness constraints applied to single fields. On multiple fields, undefined values are taken into account. That is, sets of values are considered as being duplicated if they have the same defined and undefined values.

Additional localized validation messages can be defined using the element `osd:validation` under the elements `annotation/appinfo`. If no custom validation messages are defined, a built-in validation message will be used.

Limitations:

1. The target of the `xs:field` element must be in a table.
2. The uniqueness constraint does not apply to fields inside an aggregated list.
3. The uniqueness constraint does not apply to computed fields.

See also *Uniqueness constraint in the Java API* `UniquenessConstraintAPI`

## 83.2 Extended facets

EBX® provides additional constraints that are not specified in XML Schema, but that are useful for managing master data.

In order to guarantee XML Schema conformance, these extended facets are defined under the element `annotation/appinfo/otherFacets`.

### Foreign keys

EBX® allows to create a reference to an existing table by means of a specific facet. See [Foreign keys](#) [p 500] for more information.

### Dynamic constraints

Dynamic constraint facets retain the semantics of XML Schema, but the value attribute is replaced with a path attribute that allows fetching the value from another element. The available dynamic constraints are:

- length

- minLength
- maxLength
- maxInclusive
- maxExclusive
- minInclusive
- minExclusive

Using these facets, the data model can be modified dynamically.

### Example:

```
<xs:element name="amount">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:minInclusive path="/domain/Loan/Pricing/AmountMini/amount" />
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:element>
```

In this example, the boundary of the facet `minInclusive` is not statically defined. The value of the boundary comes from the node `/domain/Loan/Pricing/AmountMini/amount`.

### Restrictions:

- Target field cannot be an aggregated list. That is, it cannot define `maxOccurs = 1`.
- Data type of the target field must be compatible with the facet. That is, it must be:
  - of type integer for facets `length`, `minLength` and `maxLength`.
  - compatible with the data type of the field holding the facet for facets `maxInclusive`, `maxExclusive`, `minInclusive` and `minExclusive`.
- Target field cannot be in a table if the field holding the facet is not in a table.
- Target field must be in the same table or outside a table if the field holding the facet is in a table.
- If the target field is under one or more aggregated lists, the field holding the facet must also be under these aggregated lists. That is: the field holding the facet must be in the same list occurrence as the target field, or in a parent occurrence, so that the target field refers to a single value, from an XPath perspective.

## FacetOResource constraint

This facet must be defined for every definition using the type `osd:resource`, to specify the subset of available packaged resource files as an enumeration. For more information on this type, see [osd:resource type](#) [p 486]. It has the following attributes:

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>moduleName</b>   | Indicates, using an alias, the EBX® module that contains the resource. If the resource is contained in the current module, the alias must be preceded by "wbp". Otherwise, the alias must be one of the values defined in the element <code>&lt;dependencies&gt;</code> in the file <code>module.xml</code> .                                                                                                                                                                                                                                                                                                                                      |
| <b>resourceType</b> | Represents the resource type that is one of the following values: 'Image', 'JavaScript', 'Style sheet', 'HTML'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>relativePath</b> | Indicates in which directory the resources will be located. This directory must be located under the directory that corresponds to the resource type. For example, for an "Image" type resource, the directory <code>www/common/images/</code> , located at the same level as the directory <code>WEB-INF/</code> of the target module, will be used and the relative path will have to be defined from this. Furthermore, if a resource is defined in a localized directory ( <code>www/fr/</code> for example), it will only be taken into account if another resource with the same name is defined in the directory <code>www/common/</code> . |

This facet has the same behavior as an enumeration facet: the values are collected by recursively listing all the files in the local path in the specified resource type directory in the specified module.

### Example:

```
<xs:element name="promotion" type="osd:resource">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:FacetOResource osd:moduleName="wbp"
          osd:resourceType="ext-images" osd:relativePath="promotion/" />
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

For an overview of the standard directory structure of an EBX® module (Java EE web application), see [Module structure](#) [p 461].

## excludeValue constraint

This facet verifies that a value is not the same as the specified excluded value.

In this example, the empty string is excluded from the allowed values.

### Example:

```
<xs:element name="roleName">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:excludeValue value="">
          <osd:validation>
```



```

    <severity>error</severity>
    <message>Please select address role(s).</message>
  </osd:validation>
</osd:excludeValue>
</osd:otherFacets>
</xs:appinfo>
</xs:annotation>
<xs:simpleType type="xs:string" />
</xs:element>

```

## ***excludeSegment constraint***

This facet verifies that a value is not included in a range of values. Boundaries are excluded.

### **Example:**

In this example, values between 20000 and 20999 are not allowed.

```

<xs:element name="zipCode">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:excludeSegment minValue="20000" maxValue="20999">
          <osd:validation>
            <severity>error</severity>
            <message>Postal code not valid.</message>
          </osd:validation>
        </osd:excludeSegment>
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
  <xs:simpleType type="xs:string" />
</xs:element>

```

## ***Enumeration facet defined using another node***

By default, an enumeration facet is described statically in XML Schema.

The content of an enumeration facet can also be provided dynamically by a list of simple elements in the data model.

### **Example:**

In this example, the content of an enumeration facet is sourced from the node CountryList.

```

<xs:annotation>
  <xs:appinfo>
    <osd:otherFacets>
      <osd:enumeration osd:path="../../CountryList" />
    </osd:otherFacets>
  </xs:appinfo>
</xs:annotation>

```

The referred node CountryList:

- Must be an aggregated list, that is, `maxOccurs > 1`.
- Must be a list of elements of the same type as the node with the enumeration facet.
- Must be a node outside a table if the node with the enumeration facet is not inside a table.
- Must be a node outside a table or in the same table as the node with the enumeration facet if the node with this enumeration is inside a table.
- If the target field is under one or more aggregated lists, the field holding the facet must also be under these aggregated lists. That is: the field holding the facet must be in the same list occurrence as the target field, or in a parent occurrence, so that the target field refers to a single value, from an XPath perspective.

### **Example:**

```

<xs:element name="FacetEnumBasedOnList">
  <xs:complexType>

```

```

<xs:sequence>
  <xs:element name="CountryList" maxOccurs="unbounded">
    <xs:simpleType>
      <xs:restriction base="xs:string">
        <xs:enumeration value="DE" osd:label="Germany" />
        <xs:enumeration value="AT" osd:label="Austria" />
        <xs:enumeration value="BE" osd:label="Belgium" />
        <xs:enumeration value="JP" osd:label="Japan" />
        <xs:enumeration value="KR" osd:label="Korea" />
        <xs:enumeration value="CN" osd:label="China" />
      </xs:restriction>
    </xs:simpleType>
  </xs:element>
  <xs:element name="CountryChoice" type="xs:string">
    <xs:annotation>
      <xs:appinfo>
        <osd:otherFacets>
          <osd:enumeration osd:path="../../CountryList" />
        </osd:otherFacets>
      </xs:appinfo>
    </xs:annotation>
  </xs:element>
</xs:sequence>
</xs:complexType>
</xs:element>

```

## 83.3 Programmatic facets

A programmatic constraint can be added to any XML element declaration for a simple type.

In order to guarantee XML Schema conformance, programmatic constraints are specified under the element annotation/appinfo/otherFacets.

### ***Programmatic constraints***

A programmatic constraint is defined by a Java class that implements the interface `Constraint`<sup>API</sup>.

As additional parameters can be defined, the implemented Java class must conform to the JavaBean protocol.

#### **Example:**

In the example below, the Java class must define the methods: `getParam1()`, `setParam1(String)`, `getParamX()`, `setParamX(String)`, etc.

```

<xs:element name="amount">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:constraint class="com.foo.CheckAmount">
          <param1>...</param1>
          <param...n>...</param...n>
        </osd:constraint>
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:element>

```

See also *JavaBean specifications* Package `com.orchestranetworks.schema.JavaBeans`<sup>API</sup>

### ***Programmatic enumeration constraints***

An enumeration constraint adds an ordered list of values to a basic programmatic constraint. This facet allows selecting a value from a list. It is defined by a Java class that implements the interface `ConstraintEnumeration`<sup>API</sup>.

#### **Example:**

```

<xs:element name="amount">
  <xs:annotation>

```

```

<xs:appinfo>
  <osd:otherFacets>
    <osd:constraintEnumeration class="com.foo.CheckAmountInEnumeration">
      <param1>...</param1>
      <param...n>...</param...n>
    </osd:constraintEnumeration>
  </osd:otherFacets>
</xs:appinfo>
</xs:annotation>
...
</xs:element>

```

## Constraint on 'null' values

In some cases, a value is only mandatory if some conditions are satisfied, for example, if another field has a given value. In this case, the standard XML Schema attribute `minOccurs` is insufficient because it is static.

In order to check if a value is mandatory according to its context, the following requirements must be satisfied:

1. A programmatic constraint must be defined by a Java class (see above).
2. This class must implement the interface `ConstraintOnNull`<sup>API</sup>.
3. The XML Schema cardinality attributes must specify that the element is optional (`minOccurs="0"` and `maxOccurs="1"`).

### Note

By default, constraints on 'null' values are not checked upon user input. In order to enable a check at the input, the ['checkNullInput' property](#) [p 527] must be set. Also, if the element is terminal, the dataset must also be activated.

### Example:

```

<xs:element name="amount" minOccurs="0" maxOccurs="1">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:constraint class="com.foo.CheckIfNull">
          <param1>...</param1>
          <param...n>...</param...n>
        </osd:constraint>
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:element>

```

## Constraints on table

A constraint on table is defined by a Java class that implements the interface `ConstraintOnTable`<sup>API</sup>. It can only be defined on table nodes.

As additional parameters can be defined, the implemented Java class must conform to the JavaBean protocol.

### Example:

In the example below, the Java class must define the methods: `getParam1()`, `setParam1(String)`, `getParamX()`, `setParamX(String)`, etc.

```

<xs:element name="myTable" type="MyTableType" minOccurs="0" maxOccurs="unbounded">
  <xs:annotation>
    <xs:appinfo>
      <osd:table>
        <primaryKeys>/key</primaryKeys>
      </osd:table>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:element>

```

```

<osd:otherFacets>
  <osd:constraint class="com.foo.checkTable">
    <param1>...</param1>
    <param...n>...</param...n>
  </osd:constraint>
</osd:otherFacets>
</xs:appinfo>
</xs:annotation>
</xs:element>

```

### Attention

For performance reasons, constraints on tables are only checked when getting the validation report of a dataset or table. This means that these constraints are not checked when updates, such as record insertions, deletions or modifications, occur on tables. However, the internal incremental validation framework will optimize the validation cost of these constraints if dependencies are defined. For more information, see [Validation](#) [p 298].

See also *JavaBeans Specifications* Package `com.orchestranetworks.schema.JavaBeansAPI`

## 83.4 Control policy

### ***Blocking and non-blocking constraints***

When an update in the repository is performed, and this update adds a validation error according to a given constraint, it is possible to specify whether the new error blocks the update (and cancels the transaction) or if it is considered as non-blocking (so that the update can be committed and the error

can be corrected later). The element `blocksCommit` within the element `osd:validation` allows this specification, with the following supported values:

|                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b><code>onInsertUpdateOrDelete</code></b>             | <p>Specifies that the constraint must always remain valid after an operation (dataset update, dataset deletion, record creation, update or deletion). In this case, any operation that would violate the constraint is rejected and the values remain unchanged.</p> <p>This is the default and mandatory policy for primary key constraints, data type conversion constraints (an integer or a date must be well-written) and also structural constraints in mapped tables.</p> <p>This is also the default policy for foreign key constraints that are automatically set in blocking mode (because of a table in relational mode involved in the relationship) and that do not define a control policy.</p> |
| <b><code>onUserSubmit - checkModifiedValues</code></b> | <p>Specifies that the constraint must remain valid whenever a user modifies the associated value and submits a form. In this case, any form input that would violate the constraint is rejected and the values remain unchanged.</p> <p>This is the default policy for all blocking constraints mentioned in the previous case. For example, a foreign key constraint is by default not blocking (a record referred to by other records can be deleted, etc.), except in the context of a form submit.</p>                                                                                                                                                                                                    |
| <b><code>never</code></b>                              | <p>Specifies that the constraint must never block operations. In this case, any operation that would violate the constraint is allowed. In the context of the user interface, this constraint does not block the form submission if the user sets a value that violates this constraint.</p>                                                                                                                                                                                                                                                                                                                                                                                                                  |

On foreign key constraints, the control policy that blocks all operations does not apply to filtered records. That is, a foreign key constraint is not blocking if a referenced record exists but does not satisfy a foreign key filter. In this case, updates are not rejected and a validation error occurs.

It is not possible to specify a control policy on structural constraints that are defined on relational data models or in mapped tables. That is, this property is not available for fixed length, maximum length, maximum number of digits, and decimal place constraints due to the validation policy of the underlying RDBMS blocking constraints.

This property does not apply to archive imports and when merging dataspace. That is, all blocking constraints, except structural constraints, are always disabled when importing archives and merging dataspace.

See also

[Facet validation message with severity](#) [p 536]

[Foreign keys](#) [p 500]

[Relational mode](#) [p 243]**XML Schema facet**

The control policy is described by the element `osd:validation` in `annotation/appinfo` under the definition of the facet.

**Example:**

```
<xs:element name="zipCode">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:minInclusive value="1000">
        <xs:annotation>
          <xs:appinfo>
            <osd:validation>
              <blocksCommit>onInsertUpdateOrDelete</blocksCommit>
            </osd:validation>
          </xs:appinfo>
        </xs:annotation>
      </xs:minInclusive>
    </xs:restriction>
  </xs:simpleType>
</xs:element>
```

**XML Schema enumeration facet**

The control policy is described by the element `osd:enumerationValidation` in `annotation/appinfo` under the definition of the field.

**Example:**

```
<xs:element name="Gender">
  <xs:annotation>
    <xs:appinfo>
      <osd:enumerationValidation>
        <blocksCommit>onInsertUpdateOrDelete</blocksCommit>
      </osd:enumerationValidation>
    </xs:appinfo>
  </xs:annotation>
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="0" osd:label="male" />
      <xs:enumeration value="1" osd:label="female" />
    </xs:restriction>
  </xs:simpleType>
</xs:element>
```

**EBX® facet**

The control policy is described by the element `osd:validation` under the definition of the facet (which is defined in `annotation/appinfo/otherFacets`).

The control policy with values `onInsertUpdateOrDelete` and `onUserSubmit-checkModifiedValues` is only available on `osd:excludeSegment`, `osd:excludeValue` and `osd:tableRef` EBX® facets.

The control policy with the value `never` can be defined on all EBX® facets. On programmatic constraints, the control policy with the value `never` can only be set directly during the setup of the corresponding constraint. See `ConstraintContext.setBlocksCommitToNeverAPI` and `ConstraintContextOnTable.setBlocksCommitToNeverAPI` in the Java API for more information.

**Example:**

```
<xs:element name="price" type="xs:decimal">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:minInclusive path="../../priceMin">
          <osd:validation>
            <blocksCommit>onInsertUpdateOrDelete</blocksCommit>
          </osd:validation>
        </osd:minInclusive>
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

```
</osd:otherFacets>
</xs:appinfo>
</xs:annotation>
</xs:element>
```

### Check 'null' input

According to the EBX® default validation policy, in order to allow temporarily incomplete input, a mandatory element is not checked for completion upon user input. Rather, it is verified at the dataset validation only. If completion must be checked immediately upon user input, the element must additionally specify the attribute `osd:checkNullInput="true"`.

#### Note

A value is mandatory if the data model specifies a mandatory element, either statically, using `minOccurs="1"`, or dynamically, using a constraint on 'null'. For terminal elements, mandatory values are only checked for an activated dataset. For non-terminal elements, the dataset does not need to be activated.

#### Example:

```
<xs:element name="amount" osd:checkNullInput="true" minOccurs="1">
...
</xs:element>
```

#### See also

[Constraint on 'null'](#) [p 523]

[Whitespace management](#) [p 527]

[Empty string management](#) [p 529]

### EBX® whitespace management for data types

According to XML Schema (see <https://www.w3.org/TR/xmlschema-2/#rf-whiteSpace>), whitespace handling must follow one of the procedures *preserve*, *replace* or *collapse*:

|                 |                                                                                                                                                                           |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>preserve</b> | No normalization is performed, the value is unchanged.                                                                                                                    |
| <b>replace</b>  | All occurrences of #x9 (tab), #xA (line feed) and #xD (carriage return) are replaced with #x20 (space).                                                                   |
| <b>collapse</b> | After the processing according to the replace procedure, contiguous sequences of #x20 are then collapsed to a single #x20, and any leading or trailing #x20s are removed. |

### General whitespace handling

EBX® complies with the XML Schema recommendation:

- For fields of type `xs:string`, whether a primary key element or not, whitespaces are always preserved and an empty string is never converted to `null`.

- For other fields (non-`xs:string` type), whitespaces are always collapsed and empty strings are converted to `null`.

### Attention

Exceptions:

- For fields of type `osd:html` or `osd:password`, whitespaces are always preserved and empty strings are converted to `null`.
- For fields of type `xs:string` that define the property `osd:checkNullInput="true"`, an empty string is interpreted as `null` at user input by EBX®.

## Whitespace handling upon user input

The rules described in the previous section are applied in the user interface, but leading and trailing whitespaces are removed upon user input. That is, in the user interface, whitespaces are by default always trimmed upon user input. Other input methods (Import XML/CSV, Data services, API updates) are not trimmed from the user interface.

### Attention

Exceptions:

- For fields of type `osd:password`, whitespaces are not trimmed upon user input.
- For foreign key fields, whitespaces are not trimmed upon user input.

It is possible to indicate in a data model that whitespaces should not be trimmed upon user input. The attribute `osd:trim="disable"` can be set on the fields that allow leading and trailing whitespaces upon user input.

### Example:

```
<xs:element name="field" osd:trim="disable" type="xs:string">
  ...
</xs:element>
```

## Whitespace handling for primary keys of type string

For primary key columns of type `xs:string`, a default EBX® constraint is defined. This constraint forbids empty strings and non-collapsed whitespace values when creating a record. That is, any record creation that would violate this constraint is rejected.

However, if the primary key node specifies its own `xs:pattern` facet, this facet overrides the default EBX® constraint. For example, the specific pattern `".*"` would accept any string, although this is not recommended.

The default constraint allows handling certain ambiguities. For example, it would be difficult for a user to distinguish between the following strings: "12 34" and "12 34". For generic values, this would not create conflicts, however, errors would occur for primary keys.

See also [Tables and relationships](#) [p 495]



## Empty string management

### Default conversion

For nodes of type `xs:string`, no distinction is made at user input between an empty string and a `null` value. That is, an empty string value is automatically converted to `null` at user input.

### Distinction between empty strings and 'null' value

There are certain cases where the distinction is made between an empty string and the `null` value, such as when:

- A primary key defines a pattern that allows empty strings.
- An element defines a foreign key constraint and the target table has a single primary key defining a pattern that allows empty strings.
- An element defines a static enumeration that contains an empty string.
- An element defines a dynamic enumeration to another element with one of the aforementioned cases.

If the distinction is made between an empty string and a `null` value, this implies the following behaviors:

- An empty string will not be converted to `null` at user input,
- Input fields for nodes of type `xs:string` will display an additional button for setting the value of the node to `null`,
- At validation time, an empty string will be considered to be a compliant value with regard to the `minOccurs="1"` property.

#### Attention

In relational mode, the Oracle database does not support the distinction between empty strings and `null` values, and these specific cases are not supported.

See also [Relational mode](#) [p 243]

## 83.5 Triggers and functions

### Computed values

By default, data is read and persisted in the XML repository. Nevertheless, data may be the result of a computation and/or external database access, for example, an RDBMS or a central system.

EBX® allows taking into account other data in the current dataset context.

This is made possible by defining *functions*.

A function is specified in the data model using the `osd:function` element (see example below).

- The value of the `class` attribute must be the qualified name of a Java class that implements the Java interface `ValueFunctionAPI`.

- Additional parameters may be specified at the data model level, in which case the JavaBean convention is applied.

**Example:**

```
<xs:element name="computedValue">
  <xs:annotation>
    <xs:appinfo>
      <osd:function class="com.foo.ComputeValue">
        <param1>...</param1>
        <param...n>...</param...n>
      </osd:function>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:element>
```

In some cases, it can be useful to disable the validation of computed values if the execution of a function is time-consuming. Indeed, if the function is attached to a table with N records, then it will be called N times when validating this table. The property `osd:disableValidation= "true"` specified in the data model allows to disable the validation of a computed value (see example below).

**Example:**

```
<xs:element name="computedValue" osd:disableValidation="true">
  <xs:annotation>
    <xs:appinfo>
      <osd:function class="com.foo.ComputeValue">
        <param1>...</param1>
        <param...n>...</param...n>
      </osd:function>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:element>
```

## Triggers

Datasets or table records can be associated with methods that are automatically executed when some operations are performed, such as creations, updates, or deletions.

In the data model, these triggers must be declared under the `annotation/appinfo` element using the `osd:trigger` element.

For dataset triggers, a Java class that extends the abstract class `InstanceTriggerAPI` must be declared inside the element `osd:trigger`.

In the case of dataset triggers, it is advised to define `annotation/appinfo/osd:trigger` tags just under the root element of the data model.

**Example:**

```
<xs:element name="root" osd:access="- -">
  ...
  <xs:annotation>
    <xs:appinfo>
      <osd:trigger class="com.foo.MyInstanceTrigger">
        <param1>...</param1>
        <param...n>...</param...n>
      </osd:trigger>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:element>
```

For the definition of table record triggers, a Java class that extends the abstract class `TableTriggerAPI` must be defined inside the `osd:trigger` element. It is advised to define the `annotation/appinfo/osd:trigger` elements just under the element describing the associated table or table type.

**Examples:**

On a table element:

```
<xs:element name="myTable" type="MyTableType" minOccurs="0" maxOccurs="unbounded">
  <xs:annotation>
    <xs:appinfo>
      <osd:table>
        <primaryKeys>/key</primaryKeys>
      </osd:table>
      <osd:trigger class="com.foo.MyTableTrigger" />
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

On a table type element:

```
<xs:complexType name="MyTableType">
  ...
  <xs:annotation>
    <xs:appinfo>
      <osd:trigger class="com.foo.MyTableTrigger">
        <param1>...</param1>
        <param...n>...</param...n>
      </osd:trigger>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:complexType>
```

As additional parameters can be defined, the implemented Java class must conform to the JavaBean protocol. In the example above, the Java class must define the methods: `getParam1()`, `setParam1(String)`, `getParamX()`, `setParamX(String)`, etc.

## Auto-incremented values

It is possible to define auto-incremented values. Auto-incremented values are only allowed inside tables, and they must be of the type `xs:int` or `xs:integer`.

An auto-increment is specified in the data model using the element `osd:autoIncrement` under the element annotation/appinfo.

### Example:

```
<xs:element name="autoIncrementedValue" type="xs:int">
  <xs:annotation>
    <xs:appinfo>
      <osd:autoIncrement />
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

Also, there are two optional elements, `start` and `step`:

- The `start` attribute specifies the first value for this auto-increment. If this attribute is not specified, then the value 1 is set by default.
- The `step` attribute specifies the step for the next value to be generated by the auto-increment. If this attribute is not specified, then the value 1 is set by default.

### Example:

```
<xs:element name="autoIncrementedValue" type="xs:int">
  <xs:annotation>
    <xs:appinfo>
      <osd:autoIncrement>
        <start>100</start>
        <step>5</step>
      </osd:autoIncrement>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

A field specifying an `osd:autoIncrement` has the following behavior:

- The computation and allocation of the field value are performed whenever a new record is inserted and the field value is undefined.
- No allocation is performed if a programmatic insertion already specifies a non-null value. For example, if an archive import or an XML import specifies the value, that value is preserved. Consequently, the allocation is not performed for a record insertion in occulting or overwriting modes.
- A newly allocated value is, whenever possible, unique in the scope of the repository. More precisely, the uniqueness of the allocation spans over all the datasets of the data model, and it also spans over all the dataspace. The latter case allows the merge of a dataspace into its parent with a reasonable guarantee that there will be no conflict if the `osd:autoIncrement` is part of the records' primary key.

This principle has a very specific limitation: when a mass update transaction that specifies values is performed at the same time as a transaction that allocates a value on the same field, it is possible that the latter transaction will allocate a value that will be set by the first transaction (there is no locking between different dataspace).

Internally, the auto-increment value is stored in the 'Auto-increments' table of the repository. In the user interface, it can be accessed by administrators in the 'Administration' area. This field is automatically updated so that it defines the greatest value ever set on the associated `osd:autoIncrement` field, in any instance or dataspace in the repository. This value is computed, taking into account the max value found in the table being updated.

In certain cases, for example when multiple environments have to be managed (development, test, production), each with different auto-increment ranges, it may be required to avoid this "max value" check. This particular behavior can be achieved using the `disableMaxTableCheck` property. It is generally not recommended to enable this property unless it is absolutely necessary, as this could generate conflicts in the auto-increment values. However, this property can be set in the following ways:

- Locally, by setting a parameter element in the auto-increment declaration:  
`<disableMaxTableCheck>true</disableMaxTableCheck>`,
- For the whole data model, by setting `<osd:autoIncrement disableMaxTableCheck="true"/>` in the element `xs:appinfo` of the data model declaration, or
- Globally, by setting the property `ebx.autoIncrement.disableMaxTableCheck=true` in the EBX® main configuration file.

See [TIBCO EBX® main configuration file](#) [p 349].

#### Note

When this option is enabled globally, it becomes possible to create records in the table of auto-increments, for example by importing from XML or CSV. If this option is not selected, creating records in the table of auto-increments is prohibited to ensure the integrity of the repository.

# Labels and messages

TIBCO EBX® allows to have custom labels and error messages for data models to be displayed in the interface.

This chapter contains the following topics:

- 1. [Label and description](#)
- 2. [Enumeration labels](#)
- 3. [Mandatory error message \(osd:mandatoryErrorMessage\)](#)
- 4. [Conversion error message](#)
- 5. [Facet validation message with severity](#)

## 84.1 Label and description

A label and a description can be added to each node in an adaptation model.

In EBX®, each adaptation node is displayed with its label. If no label is defined, the name of the element is used.

Two different notations can be used:

|        |                                                                                                                                                                  |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Full   | The label and description are defined by the elements <osd:label> and <osd:description> respectively.                                                            |
| Simple | The label is extracted from the text content, ending at the first period ('.'), with a maximum of 60 characters. The description uses the remainder of the text. |

The description may also have a hyperlink, either a standard HTML href to an external document, or a link to another node of the adaptation within EBX®.

- When using the href notation or any other HTML, it must be properly escaped.
- EBX® link notation is not escaped and must specify the path of the target, for example:  
`<osd:link path="../../../misc1">Link to another node in the adaptation</osd:link>`

**Example:**

```
<xs:element name="misc1" type="xs:string">
  <xs:annotation>
    <xs:documentation>
```

```

Miscellaneous 1. This is the description of miscellaneous element #1.
Click <a href="https://www.tibco.com" target="_blank">here</a>
to learn more.
</xs:documentation>
</xs:annotation>
</xs:element>
<xs:element name="misc2" type="xs:string">
  <xs:annotation>
    <xs:documentation>
      <osd:label>
        Miscellaneous 2
      </osd:label>
      <osd:description>
        This is the miscellaneous element #2 and here is a
        <osd:link path="../misc1"> link to another node in the
        adaptation</osd:link>.
      </osd:description>
    </xs:documentation>
  </xs:annotation>
</xs:element>

```

If a node points to a named type, then the label of the node replaces the label of the named type. The same mechanism applies to the description of the node (element `osd:description`).

#### Note

Regarding whitespace management, the label of a node is always *collapsed* when displayed. That is, contiguous sequences of blanks are collapsed to a single blank, and leading and trailing blanks are removed. In descriptions, however, whitespaces are always *preserved*.

## Dynamic labels and descriptions

As an alternative to statically defining the localized labels and descriptions for each node, it is possible to specify a Java class that programmatically determines the labels and descriptions for the nodes of the data model. To define the class, include the element `osd:documentation`, with the attribute `class` in the data model. It is possible to pass JavaBean properties using nested parameter elements.

### Example:

```

<xs:schema ...>
  <xs:annotation>
    <xs:appinfo>
      <osd:documentation class="com.foo.MySchemaDocumentation">
        <param1>...</param1>
        <param2>...</param2>
      </osd:documentation>
    </xs:appinfo>
  </xs:annotation>
  ...
</xs:schema ...>

```

The labels and descriptions that are provided programmatically take precedence over the ones defined locally on individual nodes.

See also *SchemaDocumentation*<sup>API</sup>

## 84.2 Enumeration labels

In an enumeration, a simple, non-localized label can be added to each enumeration element, using the attribute `osd:label`.

### Attention

Labels defined for an enumeration element are always collapsed when displayed.

**Example:**

```
<xs:element name="Service" maxOccurs="unbounded">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="1" osd:label="Blue" />
      <xs:enumeration value="2" osd:label="Red" />
      <xs:enumeration value="3" osd:label="White" />
    </xs:restriction>
  </xs:simpleType>
</xs:element>
```

It is also possible to fully localize the labels using the standard `xs:documentation` element. If both non-localized and localized labels are added to an enumeration element, the non-localized label will be displayed in any locale that does not have a label defined.

**Example:**

```
<xs:element name="access" minOccurs="0">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="readOnly">
        <xs:annotation>
          <xs:documentation xml:lang="en-US">
            read only
          </xs:documentation>
          <xs:documentation xml:lang="fr-FR">
            lecture seule
          </xs:documentation>
        </xs:annotation>
      </xs:enumeration>
      <xs:enumeration value="readWrite">
        <xs:annotation>
          <xs:documentation xml:lang="en-US">
            read/write
          </xs:documentation>
          <xs:documentation xml:lang="fr-FR">
            lecture écriture
          </xs:documentation>
        </xs:annotation>
      </xs:enumeration>
      <xs:enumeration value="hidden">
        <xs:annotation>
          <xs:documentation xml:lang="en-US">
            hidden
          </xs:documentation>
          <xs:documentation xml:lang="fr-FR">
            masqué
          </xs:documentation>
        </xs:annotation>
      </xs:enumeration>
    </xs:restriction>
  </xs:simpleType>
</xs:element>
```

## 84.3 Mandatory error message (`osd:mandatoryErrorMessage`)

If the node specifies the attribute `minOccurs="1"` (default behavior), then an error message, which must be provided, is displayed if the user does not complete the field. This error message can be defined specifically for each node using the element `osd:mandatoryErrorMessage`.

**Example:**

```
<xs:element name="birthDate" type="xs:date">
  <xs:annotation>
    <xs:documentation>
      <osd:mandatoryErrorMessage>
        Please give your birth date.
      </osd:mandatoryErrorMessage>
    </xs:documentation>
  </xs:annotation>
</xs:element>
```

The mandatory error message can be localized:

```
<xs:documentation>
  <osd:mandatoryErrorMessage xml:lang="en-US">
    Name is mandatory
  </osd:mandatoryErrorMessage>
  <osd:mandatoryErrorMessage xml:lang="fr-FR">
    Nom est obligatoire
  </osd:mandatoryErrorMessage>
</xs:documentation>
```

#### Note

Regarding whitespace management, the enumeration labels are always *collapsed* when displayed.

## 84.4 Conversion error message

For each predefined XML Schema element, it is possible to define a specific error message if the user entry has an incorrect format.

#### Example:

```
<xs:element name="email" type="xs:string">
  <xs:annotation>
    <xs:documentation>
      <osd:ConversionErrorMessage xml:lang="en-US">
        Please enter a valid email address.
      </osd:ConversionErrorMessage>
      <osd:ConversionErrorMessage xml:lang="fr-FR">
        Saisissez un e-mail valide.
      </osd:ConversionErrorMessage>
    </xs:documentation>
  </xs:annotation>
</xs:element>
```

## 84.5 Facet validation message with severity

The validation message that is displayed when the value of a field does not comply with a constraint can define a custom severity, a default non-localized message, and localized message variants. If no severity is specified, the default level is error. Blocking constraints *must* have the severity error.

### **XML Schema facet (osd:validation)**

The validation message is described by the element `osd:validation` in `annotation/appinfo` under the definition of the facet.

#### Example:

```
<xs:element name="zipCode">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <!--facet is not localized, but validation message is localized-->
      <xs:minInclusive value="01000">
        <xs:annotation>
          <xs:appinfo>
            <osd:validation>
              <severity>error</severity>
              <message>Non-localized message.</message>
              <message xml:lang="en-US">English error message.</message>
              <message xml:lang="fr-FR">Message d'erreur en français.</message>
            </osd:validation>
          </xs:appinfo>
        </xs:annotation>
      </xs:minInclusive>
    </xs:restriction>
  </xs:simpleType>
</xs:element>
```



## ***XML Schema enumeration facet (osd:enumerationValidation)***

The validation message is described by the element `osd:enumerationValidation` in `annotation/appinfo` under the definition of the field.

### **Example:**

```
<xs:element name="Gender">
  <xs:annotation>
    <xs:appinfo>
      <osd:enumerationValidation>
        <severity>error</severity>
        <message>Non-localized message.</message>
        <message xml:lang="en-US">English error message.</message>
        <message xml:lang="fr-FR">Message d'erreur en français.</message>
      </osd:enumerationValidation>
    </xs:appinfo>
  </xs:annotation>
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="0" osd:label="male" />
      <xs:enumeration value="1" osd:label="female" />
    </xs:restriction>
  </xs:simpleType>
</xs:element>
```

## ***EBX® facet (osd:validation)***

The validation message is described by the element `osd:validation` under the definition of the facet (which is defined in `annotation/appinfo/otherFacets`).

### **Example:**

```
<xs:element name="price" type="xs:decimal">
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:minInclusive path="../priceMin">
          <osd:validation>
            <severity>error</severity>
            <message>Non-localized message.</message>
            <message xml:lang="en-US">English error message.</message>
            <message xml:lang="fr-FR">Message d'erreur en français.</message>
          </osd:validation>
        </osd:minInclusive>
      </osd:otherFacets>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```



## CHAPTER 85

# Additional properties

This chapter contains the following topics:

1. [Default values](#)
2. [Access properties](#)
3. [Information](#)
4. [Default view](#)
5. [Comparison mode](#)
6. [Apply last modifications policy](#)
7. [Categories](#)

## 85.1 Default values

In a data model, it is possible to specify a default value for a field using the attribute `default`. This property is used to assign a default value if no value is defined for a field.

The default value is displayed in the user input field at the creation time. That is, the default value will be displayed when creating a new record or adding a new occurrence to an aggregated list.

**Example:**

In this example, the element specifies a default string value.

```
<xs:element name="fieldwithDefaultValue" type="xs:string" default="aDefaultValue" />
```

## 85.2 Access properties

The attribute `osd:access` defines the access mode, that is, whether the data of a particular data model node can be read and/or written. This attribute must have one of the following values: `RW`, `R-`, `CC` or `--`.

For each XML Schema node, three types of adaptation are possible:

1. *Adaptation terminal node*

This node is displayed with an associated value in TIBCO EBX®. When accessed using the method `Adaptation.get()`, it uses the adaptation search algorithm.

2. *Adaptation non-terminal node*

This node is a complex type that is only displayed in EBX® if it has one child node that is also an adaptation terminal node. It has no value of its own. When accessed using the method `Adaptation.get()`, it returns `null`.

3. *Non-adaptable node*

This node is not an adaptation terminal node and has no child adaptation terminal nodes. This node is never displayed in EBX®. When accessing using the method `Adaptation.get()`, it returns the node default value if one is defined, otherwise it returns `null`.

See also *Adaptation<sup>API</sup>*

| Access mode | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RW          | <i>Adaptation terminal node:</i> value can be read and written in EBX®.                                                                                                                                                                                                                                                                                                                                                                                                           |
| R-          | <i>Adaptation terminal node:</i> value can only be read in EBX®.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| CC          | <i>Cut:</i> This is not an adaptation terminal node and none of its children are adaptation terminal nodes. This "instruction" has priority over any child node regardless of the value of their <i>access</i> attribute.                                                                                                                                                                                                                                                         |
| --          | If the node is a simple type, it is not adaptable. If the node is a complex type, it is not an adaptation terminal node and does not define any child nodes.<br><br>The root node of a data model must specify this access mode.                                                                                                                                                                                                                                                  |
| Default     | If the <i>access</i> attribute is not defined: <ul style="list-style-type: none"><li>• If the node is a computed value, it is considered to be R-</li><li>• If the node is a simple type and its value is not computed, it is considered to be RW</li><li>• If the node is an aggregated list, it is then a terminal value and is considered to be RW</li><li>• Otherwise, it is not an adaptation terminal node and it does not define anything about its child nodes.</li></ul> |

**Example:**

In this example, the element is adaptable because it is an adaptation terminal node.

```
<xs:element name="proxyIpAddress" type="xs:string" osd:access="RW"/>
```

# 85.3 Information

The element `osd:information` allows specifying additional information. This information can then be used by the integration code, for any purpose, by calling the method `SchemaNode.getInformationAPI`.

**Example:**

```
<xs:element name="misc" type="xs:string">
  <xs:annotation>
    <xs:appinfo>
      <osd:information>
        This is the text information of miscellaneous element.
      </osd:information>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

```
</xs:annotation>
</xs:element>
```

## 85.4 Default view

### *Hiding a field or a table in the default view*

It is possible for a table or field inside a table to be hidden by default in EBX® by using the element `osd:defaultView/hidden`. This property is used to hide elements from the default view of a dataset without defining specific access permissions. That is, elements hidden by default will not be visible in any default forms and views, whether tabular or hierarchical, generated from the structure of the associated data model.

#### **Attention**

- If an element is configured to be hidden in the default view of a dataset, then the access permissions associated with this field will not be evaluated.
- It is possible to display a field that is hidden in the default view of a dataset by defining a view. Only in this case will the access permissions associated with this field be evaluated to determine whether the field will be displayed or not.
- It is not possible to display a table that is hidden in the default view of a dataset (in the navigation pane).

#### **Example:**

In this example, the element is hidden in the default view of a dataset.

```
<xs:element name="hiddenField" type="xs:string" minOccurs="0"/>
  <xs:annotation>
    <xs:appinfo>
      <osd:defaultView>
        <hidden>true</hidden>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

### *Hiding groups and fields in views*

It is possible for a field or a group to be hidden in all views of a table by using the element `osd:defaultView/hiddenInViews`. This property is used to hide elements from the tabular (including the default tabular view) and hierarchical views of a dataset without defining specific access permissions. That is, hidden elements will not be visible in any views, whether tabular or hierarchical, created from the structure of the associated data model. However, hidden elements in views will be displayed in forms.

To specify whether or not to hide an element in all views, use the `osd:defaultView/hiddenInViews="true|false"` element.

If this property is set to `true`, then the element will not be selectable when creating a custom view. As a consequence, the element will not be displayed in all views of a table in a dataset.

If a group is configured as hidden in views, then all the fields nested under this group will not be displayed respectively in the views of the table.

## Hiding a field in search tools

To specify whether or not to hide an element in search tools, use the element `osd:defaultView/hiddenInSearch="true|false|textSearchOnly"`.

If this property is set to `true`, then the field will not be selectable in the text and typed search tools of a dataset.

If this property is set to `textSearchOnly`, then the field will not be selectable only in the text search of a dataset but will be selectable in the typed search.

### Note

If a group is configured as hidden in search tools or only in the text search, then all the fields nested under this group will not be displayed respectively in the search tools or only in the text search.

### Example:

```
<xs:element name="hiddenFieldInSearch" type="xs:string" minOccurs="0"/>
  <xs:annotation>
    <xs:appinfo>
      <osd:defaultView>
        <hiddenInSearch>true</hiddenInSearch>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

In this example, the element is hidden in the text and typed search tools of a dataset.

```
<xs:element name="hiddenFieldOnlyInTextSearch" type="xs:string" minOccurs="0"/>
  <xs:annotation>
    <xs:appinfo>
      <osd:defaultView>
        <hiddenInSearch>textSearchOnly</hiddenInSearch>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

In this example, the element is hidden only in the text search tool of a dataset.

## Hiding a field in Data Services

To specify whether or not to hide an element in data services, use the element `osd:defaultView/hiddenInDataServices`. For more information, see [Disabling fields from data model](#) [p 617].

### Note

- If a group is configured as being hidden, then all the fields nested under this group will be considered as hidden by data services.

### Example:

```
<xs:element name="hiddenFieldInDataService" type="xs:string" minOccurs="0"/>
  <xs:annotation>
    <xs:appinfo>
      <osd:defaultView>
        <hiddenInDataServices>true</hiddenInDataServices>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

In this example, the element is hidden in the Data Service select operation.

## Defining a view for the combo box selector of a foreign key

It is possible to specify a published view that will be used to display the target table or the hierarchical view of a foreign key for a smoother selection. If a view has been defined, the selector will be displayed in the user interface in the combo box of this foreign key. The definition of a view can be done by using the XML Schema element `osd:defaultView/widget/viewForAdvancedSelection`.

### Note

- This property can only be defined on foreign key fields.
- The published view must be associated with the target table of the foreign key.
- If the published view does not exist, then the advanced selection is not available in the foreign key field.

### Example:

In this example, the name of a published view is defined to display the target table of a foreign key in the advanced selection.

```
<xs:element name="catalog_ref" type="xs:string" minOccurs="0"/>
  <xs:annotation>
    <xs:appinfo>
      <osd:otherFacets>
        <osd:tableRef>
          <tablePath>/root/Catalogs</tablePath>
        </osd:tableRef>
      </osd:otherFacets>
      <osd:defaultView>
        <widget>
          <viewForAdvancedSelection>catalogView</viewForAdvancedSelection>
        </widget>
      </osd:defaultView>
    </xs:appinfo>
  </xs:annotation>
</xs:element>
```

See [Combo-box selector](#) [p 55] for more information.

## Customizing a default widget

A widget can be defined using the data model assistant. See [Default view > Widget](#) [p 55] for more information.

# 85.5 Comparison mode

The attribute `osd:comparison` can be included on a terminal node element in order to set its comparison mode. This mode controls how differences are detected for the element during comparisons. The possible values for the attribute are:

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>default</b> | Element is visible during comparisons of its data.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>ignored</b> | <p>No changes will be detected when comparing two versions of the content in records or datasets.</p> <p>During a merge, the data values of an ignored element are not merged when contents are updated, even if the values are different. For new content, the values of ignored elements are merged.</p> <p>During an archive import, values of ignored elements are not imported when contents are updated. For new content, the values of ignored elements are imported.</p> |

**Note**

- If a group is configured as being ignored during comparisons, then all the fields nested under this group will also be ignored.
- If a terminal field does not include the attribute `osd:comparison`, then it will be included by default during comparisons.

**Restrictions:**

- This property cannot be defined on non-terminal fields.
- Primary key fields cannot be ignored during comparison.

**Example:**

In this example, the first element is explicitly ignored during comparison, the second element is explicitly included.

```
<xs:element name="fieldExplicitlyIgnoredInComparison"
  type="xs:string" minOccurs="0" osd:comparison="ignored"/>
<xs:element name="fieldExplicitlyNotIgnoredInComparison"
  type="xs:string" minOccurs="0" osd:comparison="default"/>
```

# 85.6 Apply last modifications policy

The attribute `osd:applyLastModification` can be defined on a terminal node element in order to specify if this element must be included or not in the 'apply last modifications' service that can be executed in a table of a dataset.



The possible values for the attribute are:

|                |                                                                                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>default</b> | Last modifications can be applied to this element.                                                                                                                            |
| <b>ignored</b> | This element is ignored from the apply last modifications service. That is, the last modification that has been performed on this element cannot be applied to other records. |

**Note**

- If a group is configured as being ignored by the 'apply last modifications' service, then all fields nested under this group will also be ignored.
- If a terminal field does not include the attribute `osd:applyLastModification`, then it will be included by default in the apply last modifications service.

**Restriction:**

- This property cannot be defined on non-terminal fields.

**Example:**

In this example, the first element is explicitly ignored in the 'apply last modifications' service, the second element is explicitly included.

```
<xs:element name="fieldExplicitlyIgnoredInApplyLastModification"
  type="xs:string" minOccurs="0" osd:applyLastModification="ignored"/>
<xs:element name="fieldExplicitlyNotIgnoredApplyLastModification"
  type="xs:string" minOccurs="0" osd:applyLastModification="default"/>
```

# 85.7 Categories

Categories can be used for "filtering", by restricting the display of data model elements.

To create a category, add the attribute `osd:category` to a table node in the data model XSD.

## *Filters on data*

In the example below, the attribute `osd:category` is added to the node in order to create a category named *mycategory*.

```
<xs:element name="rebate" osd:category="mycategory">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="label" type="xs:string"/>
      <xs:element name="beginDate" type="xs:date"/>
      <xs:element name="endDate" type="xs:date"/>
      <xs:element name="rate" type="xs:decimal"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
```

To activate a defined category filter on a dataset in the user interface, select **Actions > Categories > <category name>** from the navigation pane.

## *Predefined categories*

Two categories with localized labels are predefined:

- Hidden

An instance node, including a table node itself, is hidden in the default view, but can be revealed by selecting **Actions > Categories > [hidden nodes]** from the navigation pane.

A table record node is always hidden.

- Constraint (deprecated)

### ***Restriction***

Categories do not apply to table record nodes, except the category 'Hidden'.

## CHAPTER 86

# Data services

This chapter details how WSDL operations' names related to a table are defined and managed by TIBCO EBX®.

This chapter contains the following topics:

1. [Definition](#)
2. [Configuration](#)
3. [Publication](#)
4. [WSDL and table operations](#)
5. [Limitations](#)

## 86.1 Definition

EBX® generates a WSDL that complies with the [W3C Web Services Description Language 1.1](#) standard. By default, WSDL operations refer to a table using the last element of the table path. A WSDL operation name is composed of the action name (prefix) and the table name (suffix). It is possible to refer to tables in WSDL operations using unique names instead of the last element of their paths by overriding the suffix operations' names.

See also [Data services using the Data Model Assistant](#) [p 41]

## 86.2 Configuration

### ***Embedded data model***

WSDL suffix operations' names are embedded in EBX®'s repository and linked to a publication. That is, when publishing an embedded data model, the list of WSDL suffix operations' names can be defined in the data model definition, under the 'Configuration > Data services' table and managed by EBX®.

### ***Packaged data model***

WSDL suffix operations' names are defined in a dedicated XML document file and must be named as the data model and end with the keyword `_entities`. For instance, if a data model is named `catalog.xsd`, then the XML document containing the configuration of the WSDL operations' names overridden will be named `catalog_entities.xml`. This XML document must also be located in the same location as the

data model. The XML document is automatically loaded by EBX® if a file that matches this pattern is found when compiling a data model.

## 86.3 Publication

The suffix operations' names are validated at compilation time and contain a list of couples containing Path with a unique table name. Checked validation rules are:

- The path is not unique,
- The table name contains a syntax error,
- The table name is not unique in the XML document.

## 86.4 WSDL and table operations

### ***WSDL Generator***

An additional validation rule has been added: a unicity check is systematically applied to table names. The SOAP operation name is composed of the operation type as a prefix and, by default, of the table name (last step of the table path) as a suffix. A dataset can contain several identical table names but with different paths. It is possible to override table names that are not unique in order to guarantee the unicity.

### ***SOAP operations***

When an operation request on table has been invoked from the SOAP connector, the target table is retrieved by priority, the name corresponds to:

1. an overridden table name,
2. the last step of the table path.

See also [Data services](#) [p 598]

## 86.5 Limitations

WSDL operations' names are not available with external data models.

# CHAPTER 87

---

## Toolbars

This chapter details how toolbars are defined and managed by TIBCO EBX®.

This chapter contains the following topics:

1. [Definition](#)
2. [Using toolbars](#)

### 87.1 Definition

Toolbars allow to customize the buttons and menus to display when accessing a table view, a hierarchical view, or a record form.

Toolbars can only be created and published using the *Data Model Assistant* and are available only on embedded and packaged data models.

For embedded data models, toolbars are embedded in EBX®'s repository and linked to a publication. That is, when publishing an embedded data model, the toolbars defined in the data model are embedded with the publication of the data model and managed by EBX®.

For packaged data models, toolbars are defined in a dedicated XML document and must be named as the data model and end with the keyword `_toolbars`. For instance, if a data model is named `catalog.xsd` then the XML document containing the definition of the toolbars must be named `catalog_toolbars.xml`. This XML document must also be placed in the same location as the data model. The toolbar document is automatically loaded by EBX® if a file complying with this pattern is found when compiling a data model.

#### See also

[Configuring toolbars using the Data Model Assistant](#) [p 73]

[Using toolbars in data models](#) [p 499]

**Toolbar API** `ToolbarFactory`<sup>API</sup>

### 87.2 Using toolbars

Toolbars can be used on tables and associations.

On tables, it is possible to specify the toolbar to display:

- On the top of a tabular view
- On each row of a tabular view

- On the top of a record form
- On the top of a hierarchical view.

On associations, it is possible to specify the toolbar to display:

- On top of the tabular view of the association
- On each row of the tabular view of the association

#### See also

[\*Using toolbars\*](#) [p 499]

[\*Associations\*](#) [p 503]

## CHAPTER 88

# Workflow model

The workflow offers two types of steps: 'library' or 'specific'.

'Library' is a bean defined in `module.xml` and is reusable. Using the 'library' bean improves the ergonomics: parameters are dynamically displayed in the definition screens.

A 'specific' object is a bean defined only by its class name. In this case, the display is not dynamic.

This chapter contains the following topics:

1. [Bean categories](#)
2. [Sample of ScriptTask](#)
3. [Sample of ScriptTaskBean](#)
4. [Samples of UserTask](#)
5. [Samples of Condition](#)
6. [Sample of ConditionBean](#)
7. [Sample of SubWorkflowsInvocationBean](#)
8. [Sample of WaitTaskBean](#)
9. [Sample of ActionPermissionsOnWorkflow](#)
10. [Sample of WorkflowTriggerBean](#)
11. [Sample of trigger starting a process instance](#)

## 88.1 Bean categories

| Step       | Library        | Specific   |
|------------|----------------|------------|
| Scripts    | ScriptTaskBean | ScriptTask |
| Conditions | ConditionBean  | Condition  |
| User task  | UserTask       |            |

## 88.2 Sample of ScriptTask

### Java Code

A script task has to override the method execute as in the following example:

```
public class NppScriptTask_CreateWorkingBranch extends ScriptTask
{
    public void executeScript(ScriptTaskContext aContext) throws OperationException
    {
        Repository repository = aContext.getRepository();
        String initialBranchString = aContext.getVariableString("initialBranch");
        AdaptationHome initialBranch = repository.lookupHome(HomeKey.forBranchName(initialBranchString));
        if (initialBranch == null)
            throw OperationException.createError("Null value for initialBranch");

        HomeCreationSpec spec = new HomeCreationSpec();
        spec.setParent(initialBranch);
        spec.setKey(HomeKey.forBranchName("Name"));
        spec.setOwner(Profile.EVERYONE);
        spec.setHomeToCopyPermissionsFrom(initialBranch);
        AdaptationHome newHome = repository.createHome(spec, aContext.getSession());
        //feeds dataContext
        aContext.setVariableString("workingBranch", newHome.getKey().getName());
    }
}
```

See also *com.orchestranetworks.workflow.ScriptTask* *ScriptTask*<sup>API</sup>

## 88.3 Sample of ScriptTaskBean

### Java Code

A script task bean has to override the method executeScript as in the following example:

```
public class ScriptTaskBean_CreateBranch extends ScriptTaskBean
{
    private String initialBranchName;

    private String newBranch;

    public String getInitialBranchName()
    {
        return this.initialBranchName;
    }

    public void setInitialBranchName(String initialBranchName)
    {
        this.initialBranchName = initialBranchName;
    }

    public String getNewBranch()
    {
        return this.newBranch;
    }

    public void setNewBranch(String newBranch)
    {
        this.newBranch = newBranch;
    }

    public void executeScript(ScriptTaskBeanContext aContext) throws OperationException
    {
        final Repository repository = aContext.getRepository();

        String initialBranchName = this.getInitialBranchName();
        final AdaptationHome initialBranch = repository.lookupHome(HomeKey.forBranchName(initialBranchName));
        final HomeCreationSpec spec = new HomeCreationSpec();
        spec.setParent(initialBranch);
        spec.setKey(HomeKey.forBranchName(XsFormats.SINGLETON.formatDateTime(new Date())));
        spec.setOwner(Profile.EVERYONE);
        spec.setHomeToCopyPermissionsFrom(initialBranch);
        final AdaptationHome branchCreate = repository.createHome(spec, aContext.getSession());
    }
}
```



```

        this.setNewBranch(branchCreate.getKey().getName());
    }
}

```

See also *com.orchestranetworks.workflow.ScriptTaskBean* *ScriptTaskBean*<sup>API</sup>

## Configuration through module.xml

A script task bean must be declared in module.xml:

```

<module>
  <beans>
    <bean className="com.orchestranetworks.workflow.genericScriptTask.ScriptTaskBean_CreateBranch">
      <documentation xml:lang="fr-FR">
        <label>Créer une branche</label>
        <description>
          Ce script permet de créer une branche
        </description>
      </documentation>
      <documentation xml:lang="en-US">
        <label>Create a branch</label>
        <description>
          This script creates a branch
        </description>
      </documentation>
      <properties>
        <property name="initialBranchName" input="true">
          <documentation xml:lang="fr-FR">
            <label>Branche initiale</label>
            <description>
              Nom de la branche initiale.
            </description>
          </documentation>
          <documentation xml:lang="en-US">
            <label>Initial branch</label>
            <description>
              Initial branch name.
            </description>
          </documentation>
        </property>
        <property name="newBranch" output="true">
          <documentation xml:lang="fr-FR">
            <label>Nouvelle branche</label>
            <description>
              Nom de la branche créée
            </description>
          </documentation>
          <documentation xml:lang="en-US">
            <label>New branch</label>
            <description>
              Created branch name.
            </description>
          </documentation>
        </property>
      </properties>
    </bean>
  </beans>
</module>

```

## 88.4 Samples of UserTask

### Service declaration via module.xml

A built-in service can be declared in module.xml to be used in the user task definition.

```

<services>
  <service name="ServiceModule">
    <resourcePath>/service.jsp</resourcePath>
    <type>branch</type>
    <documentation xml:lang="fr-FR">
      <label>Workflow service</label>
      <description>
        Ce service permet de ...
      </description>
    </documentation>
    <documentation xml:lang="en-US">

```

```

        <label>Service workflow</label>
        <description>
            The purpose of this service is ...
        </description>
    </documentation>
    <properties>
        <property name="param1" input="true">
            <documentation xml:lang="fr-FR">
                <label>Param1</label>
                <description>Param1 ...</description>
            </documentation>
        </property>
        <property name="param2" output="true">
        </property>
    </properties>
</service>
<serviceLink serviceName="adaptationService">
    <importFromSchema>
        /WEB-INF/ebx/schema/schema.xsd
    </importFromSchema>
</serviceLink>
</services>

```

## A more complex UserTask

The GUI is quite similar as the example above. The field 'Rule' must be filled to define the class extending the 'UserTask' to invoke.

```

public class NppUserTask_ValidateProduct extends UserTask
{
    public void handleWorkItemCompletion(UserTaskWorkItemCompletionContext context)
        throws OperationException
    {
        if (context.getCompletedWorkItem().isRejected())
        {
            context.setVariableString(NppConstants.VAR_VALIDATION, "KO");
            context.completeUserTask();
        }
        else if (context.checkAllWorkItemMatchStrategy())
        {
            context.setVariableString(NppConstants.VAR_VALIDATION, "OK");
            context.completeUserTask();
        }
    }

    public void handleCreate(UserTaskCreationContext context) throws OperationException
    {
        CreationWorkItemSpec spec = CreationWorkItemSpec.forOfferring(NppConstants.ROLE_PVALIDATOR);
        spec.setNotificationMail("1");
        context.createWorkItem(spec);
        context.setVariableString(NppConstants.VAR_VALIDATION, "validating");
    }
}

```

See also *com.orchestranetworks.workflow.UserTask* <sup>API</sup>

## 88.5 Samples of Condition

### Java Code

The method evaluate has to be overridden:

```

public class NppCondition_IsValidationOK extends Condition
{
    public boolean evaluateCondition(ConditionContext context) throws OperationException
    {
        String validation = context.getVariableString("validationResult");
        boolean hasError = "KO".equals(validation);
        return !hasError;
    }
}

```

See also *com.orchestranetworks.workflow.Condition* <sup>API</sup>

## 88.6 Sample of ConditionBean

### Java Code

The method `evaluateCondition` has to be overridden as in the following sample:

```
public class ConditionBean_IsBranchValid extends ConditionBean
{
    private String branchName;

    public String getBranchName()
    {
        return this.branchName;
    }

    public void setBranchName(String branchName)
    {
        this.branchName = branchName;
    }

    public boolean evaluateCondition(ConditionBeanContext aContext) throws OperationException
    {
        final Repository repository = aContext.getRepository();
        Severity severityForValidation = Severity.ERROR;
        String branchToTestName = this.getBranchName();
        final AdaptationHome branchToTest = repository.lookupHome(HomeKey.forBranchName(branchToTestName));
        if (branchToTest.getValidationReportsMap(severityForValidation) != null
            && branchToTest.getValidationReportsMap(severityForValidation).size() > 0)
        {
            return false;
        }
        return true;
    }
}
```

See also *com.orchestranetworks.workflow.ConditionBean* <sup>API</sup>

### Configuration through module.xml

The condition bean must be declared in `module.xml`:

```
<module>
  <beans>
    <bean className="com.orchestranetworks.workflow.genericScriptTask.ConditionBean_IsBranchValid">
      <documentation xml:lang="fr-FR">
        <label>Branche valide ?</label>
        <description>
          Ce script permet de tester si une branche est valide.
        </description>
      </documentation>
      <documentation xml:lang="en-US">
        <label>Branch valid ?</label>
        <description>
          This script allows to check if a branch is valid.
        </description>
      </documentation>
      <properties>
        <property name="branchName" input="true">
          <documentation xml:lang="fr-FR">
            <label>Branche à contrôler</label>
            <description>
              Nom de la branche à valider.
            </description>
          </documentation>
          <documentation xml:lang="en-US">
            <label>Branch to check</label>
            <description>
              Branch name to check.
            </description>
          </documentation>
        </property>
      </properties>
    </bean>
  </beans>
</module>
```

```
</beans>
</module>
```

## 88.7 Sample of SubWorkflowsInvocationBean

### Java Code

```
public class MySubWorkflowsInvocationBean extends SubWorkflowsInvocationBean
{
    @Override
    public void handleCreateSubWorkflows(SubWorkflowsCreationContext aContext)
        throws OperationException
    {
        final ProcessLauncher subWorkflow1 = aContext.registerSubWorkflow(
            AdaptationName.forName("validateProduct"),
            "validateProduct1");
        subWorkflow1.setLabel(UserMessage.createInfo("Validate the new product by marketing"));
        subWorkflow1.setInputParameter("workingBranch", aContext.getVariableString("workingBranch"));
        subWorkflow1.setInputParameter("code", aContext.getVariableString("code"));
        subWorkflow1.setInputParameter("service", aContext.getVariableString("marketing"));

        final ProcessLauncher subWorkflow2 = aContext.registerSubWorkflow(
            AdaptationName.forName("validateProduct"),
            "validateProduct2");
        subWorkflow2.setLabel(UserMessage.createInfo("Validate the new product by direction"));
        subWorkflow2.setInputParameter("workingBranch", aContext.getVariableString("workingBranch"));
        subWorkflow2.setInputParameter("code", aContext.getVariableString("code"));
        subWorkflow2.setInputParameter("service", aContext.getVariableString("direction"));

        // Conditional launching.
        if (aContext.getVariableString("productType").equals("book"))
        {
            final ProcessLauncher subWorkflow3 = aContext.registerSubWorkflow(
                AdaptationName.forName("generateISBN"),
                "generateISBN");
            subWorkflow3.setLabel(UserMessage.createInfo("Generate ISBN"));
            subWorkflow3.setInputParameter(
                "workingBranch",
                aContext.getVariableString("workingBranch"));
            subWorkflow3.setInputParameter("code", aContext.getVariableString("code"));
        }

        aContext.launchSubWorkflows();
    }
    @Override
    public void handleCompleteAllSubWorkflows(SubWorkflowsCompletionContext aContext)
        throws OperationException
    {
        aContext.getCompletedSubWorkflows();
        final ProcessInstance validateProductMarketing = aContext.getCompletedSubWorkflow("validateProduct1");
        final ProcessInstance validateProductDirection = aContext.getCompletedSubWorkflow("validateProduct2");
        if (aContext.getVariableString("productType").equals("book"))
        {
            final ProcessInstance generateISBN = aContext.getCompletedSubWorkflow("generateISBN");
            aContext.setVariableString("isbn", generateISBN.getDataContext().getVariableString(
                "newCode"));
        }

        if (validateProductMarketing.getDataContext().getVariableString("Accepted").equals("true")
            && validateProductDirection.getDataContext().getVariableString("Accepted").equals(
                "true"))
            aContext.setVariableString("validation", "ok");
    }
}
```

See also [\*com.orchestranetworks.workflow.SubWorkflowsInvocationBean\*](#)  
*SubWorkflowsInvocationBean*<sup>API</sup>

### Configuration through module.xml

SubWorkflowsInvocationBean bean must be declared in module.xml:

```
<module>
  <beans>
    <bean className="com.orchestranetworks.workflow.test.MySubWorkflowsInvocationBean"/>
  </beans>
```

```
</module>
```

## 88.8 Sample of WaitTaskBean

### Java Code

```
public class MyWaitTaskBean extends WaitTaskBean
{
    @Override
    public void onStart(WaitTaskOnStartContext aContext)
    {
        Map<String, String> params = new HashMap<String, String>();
        params.put("resumeId", aContext.getResumeId());
        myMethod.callWebService(params);
    }

    @Override
    public void onResume(WaitTaskOnResumeContext aContext) throws OperationException
    {
        // Defines a specific mapping.
        aContext.setVariableString("code", aContext.getOutputParameters().get("isbn"));
        aContext.setVariableString("comment", aContext.getOutputParameters().get("isbnComment"));
    }
}
```

See also *com.orchestranetworks.workflow.WaitTaskBean* WaitTaskBean<sup>API</sup>

### Configuration through module.xml

WaitTaskBean bean must be declared in module.xml:

```
<module>
  <beans>
    <bean className="com.orchestranetworks.workflow.test.MyWaitTaskBean"/>
  </beans>
</module>
```

## 88.9 Sample of ActionPermissionsOnWorkflow

### Java Code

```
package com.orchestranetworks.workflow.test;

import com.orchestranetworks.service.*;
import com.orchestranetworks.workflow.*;
import com.orchestranetworks.workflow.ProcessExecutionContext.*;

/**
 */
public class MyDynamicPermissions extends ActionPermissionsOnWorkflow
{
    public ActionPermission getActionPermission(
        WorkflowPermission aWorkflowAction,
        ActionPermissionsOnWorkflowContext aContext)
    {
        if (WorkflowPermission.VIEW.equals(aWorkflowAction)
            || WorkflowPermission.CREATE_PROCESS.equals(aWorkflowAction))
            return ActionPermission.getEnabled();
        return ActionPermission.getDisabled();
    }
}
```

See also *com.orchestranetworks.workflow.ActionPermissionsOnWorkflow* ActionPermissionsOnWorkflow<sup>API</sup>

## Configuration through module.xml

ActionPermissionsOnWorkflow bean must be declared in module.xml:

```
<module>
  <beans>
    <bean className="com.orchestranetworks.workflow.test.MyDynamicPermissions"/>
  </beans>
</module>
```

## 88.10 Sample of WorkflowTriggerBean

### Java Code

```
public class MyWorkflowTriggerBean extends WorkflowTriggerBean
{
    @Override
    public void handleAfterProcessInstanceStart(
        WorkflowTriggerAfterProcessInstanceStartContext aContext) throws OperationException
    {
        final DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());
        final MailSpec spec = aContext.createMailSpec();
        spec.notify(NotificationType.TO, "supervisor@mail.com");

        spec.setSubject("[TRIGGER] After process instance start");
        spec.setBody("The workflow '"
            + policy.formatUserMessage(aContext.getProcessInstance().getLabel())
            + "' has been created.");

        spec.sendMail(Locale.US);
    }

    @Override
    public void handleBeforeProcessInstanceTermination(
        WorkflowTriggerBeforeProcessInstanceTerminationContext aContext) throws OperationException
    {
        final DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());

        final MailSpec spec = aContext.createMailSpec();
        spec.notify(NotificationType.TO, "supervisor@mail.com");

        spec.setSubject("[TRIGGER] Before process instance termination");
        spec.setBody("The workflow '"
            + policy.formatUserMessage(aContext.getProcessInstance().getLabel())
            + "' has been completed. The created product is: '"
            + aContext.getVariableString(NppConstants.VAR_CODE) + "'.");

        spec.sendMail(Locale.US);
    }

    @Override
    public void handleAfterWorkItemCreation(WorkflowTriggerAfterWorkItemCreationContext aContext)
        throws OperationException
    {
        DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());

        MailSpec spec = aContext.createMailSpec();
        spec.notify(NotificationType.TO, "supervisor@mail.com");

        spec.setSubject("[TRIGGER] After work item creation");
        WorkItem workItem = aContext.getWorkItem();
        State state = workItem.getState();
        String body = "The work item '" + policy.formatUserMessage(workItem.getLabel())
            + "' has been created. \n The step id is : " + aContext.getCurrentStepId()
            + ". \n The work item is in state : " + policy.formatUserMessage(state.getLabel());

        if (workItem.getOfferedTo() != null)
            body += "\n The role is : " + workItem.getOfferedTo().format();
        if (workItem.getUserReference() != null)
            body += "\n The user is : " + workItem.getUserReference().format();

        spec.setBody(body);

        spec.sendMail(Locale.US);
    }

    @Override
```

```

public void handleBeforeWorkItemStart(WorkflowTriggerBeforeWorkItemStartContext aContext)
    throws OperationException
{
    DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());

    MailSpec spec = aContext.createMailSpec();
    spec.notify(NotificationType.TO, "supervisor@mail.com");

    spec.setSubject("[TRIGGER] Before work item start");
    spec.setBody("The work item '"
        + policy.formatUserMessage(aContext.getWorkItem().getLabel())
        + "' has been started. \n The current step id is : "
        + aContext.getCurrentStepId()
        + ". \n The work item user is: '"
        + DirectoryHandler.getInstance(aContext.getRepository()).displayUser(
            aContext.getWorkItem().getUserReference(),
            aContext.getSession().getLocale()) + "'");

    spec.sendMail(Locale.US);
}

@Override
public void handleBeforeWorkItemAllocation(
    WorkflowTriggerBeforeWorkItemAllocationContext aContext) throws OperationException
{
    DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());

    MailSpec spec = aContext.createMailSpec();
    spec.notify(NotificationType.TO, "supervisor@mail.com");

    spec.setSubject("[TRIGGER] Before work item allocation");
    spec.setBody("The work item '"
        + policy.formatUserMessage(aContext.getWorkItem().getLabel())
        + "' has been allocated. \n The current step id is : "
        + aContext.getCurrentStepId()
        + ". \n The work item user is: '"
        + DirectoryHandler.getInstance(aContext.getRepository()).displayUser(
            aContext.getUserReference(),
            aContext.getSession().getLocale()) + "'");

    spec.sendMail(Locale.US);
}

@Override
public void handleBeforeWorkItemDeallocation(
    WorkflowTriggerBeforeWorkItemDeallocationContext aContext) throws OperationException
{
    DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());

    MailSpec spec = aContext.createMailSpec();
    spec.notify(NotificationType.TO, "supervisor@mail.com");

    spec.setSubject("[TRIGGER] Before work item deallocation");
    spec.setBody("The work item '"
        + policy.formatUserMessage(aContext.getWorkItem().getLabel())
        + "' has been deallocated. \n The current step id is : "
        + aContext.getCurrentStepId()
        + ". \n The old work item user is: '"
        + DirectoryHandler.getInstance(aContext.getRepository()).displayUser(
            aContext.getWorkItem().getUserReference(),
            aContext.getSession().getLocale()) + "'");

    spec.sendMail(Locale.US);
}

@Override
public void handleBeforeWorkItemReallocation(
    WorkflowTriggerBeforeWorkItemReallocationContext aContext) throws OperationException
{
    DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());

    MailSpec spec = aContext.createMailSpec();
    spec.notify(NotificationType.TO, "supervisor@mail.com");

    spec.setSubject("[TRIGGER] Before work item reallocation");
    spec.setBody("The work item '"
        + policy.formatUserMessage(aContext.getWorkItem().getLabel())
        + "' has been reallocated. \n The current step id is : "
        + aContext.getCurrentStepId()
        + ". \n The work item user is: '"
        + DirectoryHandler.getInstance(aContext.getRepository()).displayUser(
            aContext.getUserReference(),
            aContext.getSession().getLocale())
        + "' The old work item user is: '"
    );
}

```

```

+ DirectoryHandler.getInstance(aContext.getRepository()).displayUser(
aContext.getWorkItem().getUserReference(),
aContext.getSession().getLocale()) + "'.");

spec.sendMail(Locale.US);

}
@Override
public void handleBeforeWorkItemTermination(
WorkflowTriggerBeforeWorkItemTerminationContext aContext) throws OperationException
{
DisplayPolicy policy = DisplayPolicyFactory.getPolicyForSession(aContext.getSession());

MailSpec spec = aContext.createMailSpec();
spec.notify(NotificationType.TO, "supervisor@mail.com");

spec.setSubject("[TRIGGER] Before work item termination");
spec.setBody("The work item '"
+ policy.formatUserMessage(aContext.getWorkItem().getLabel())
+ "' has been terminated. \n The current step id is: " + aContext.getCurrentStepId()
+ ". \n The work item has been accepted ? " + aContext.isAccepted());

spec.sendMail(Locale.US);
}
}

```

See also *com.orchestranetworks.workflow.WorkflowTriggerBean* *WorkflowTriggerBean*<sup>API</sup>

## Configuration through module.xml

WorkflowTriggerBean bean must be declared in module.xml:

```

<module>
  <beans>
    <bean className="com.orchestranetworks.workflow.test.MyWorkflowTriggerBean"/>
  </beans>
</module>

```

## 88.11 Sample of trigger starting a process instance

### Sample

```

public class TriggerWorkflow extends TableTrigger
{
    public void handleAfterModify(AfterModifyOccurrenceContext aContext) throws OperationException
    {
        ValueContext currentRecord = aContext.getOccurrenceContext();
        String code = (String) currentRecord.getValue(Path.parse("/code"));

        //Get published process
        PublishedProcessKey processPublishedKey = PublishedProcessKey.forName("productProcess");
        //Defines process instance
        ProcessLauncher launcher = ProcessLauncherHelper.createLauncher(
            processPublishedKey,
            aContext.getProcedureContext());
        //initialize Data Context
        launcher.setInputParameter("code", "/root/Client[.code=\"\" + code + "\"]");
        launcher.setInputParameter("workingBranch", aContext.getAdaptationHome().getKey().getName());

        //Starts process
        launcher.launchProcess();
    }
    //...
}

```



---

# User interface

---

## CHAPTER 89

# Interface customization

The TIBCO EBX® graphical interface can be customized through various EBX® APIs.

This chapter contains the following topics:

1. [How to embed a Web Component](#)
2. [User services](#)
3. [Form layout](#)
4. [Custom widgets](#)
5. [Table filter](#)
6. [Record label](#)
7. [CSS and JavaScript](#)

## 89.1 How to embed a Web Component

EBX® can be integrated into any application that is accessible through a supported web browser, thanks to the Web Component API.

To embed all or part of EBX® in a web page, the HTML tag `<iframe>` should be used by indicating the URL to EBX®. This URL can be specified either manually or by using the `UIHttpManagerComponent` API. A single web page may include several iframes that integrate EBX®. It is then possible to create a portal made of tables, forms, hierarchical views, etc., from EBX®.

See also [Using TIBCO EBX® as a Web Component](#) [p 191]

## 89.2 User services

A user service is an extension of EBX® that provides a graphical user interface (GUI) that allows users to access specific or advanced functions.

Powerful custom user services can be developed using the same visual components and data validation mechanisms as standard EBX® user interfaces.

See also [User service overview](#) [p 565]

## 89.3 Form layout

It is possible to override the default layout of forms in the user interface by highly customizing it with the `UIForm` API. This API provides the standard input components from EBX®, which give the possibility to customize the layout of a form, while having the same components and standard behavior as record forms using the default layout.

See also

`UIForm`<sup>API</sup>

`UIFormPaneWriter`<sup>API</sup>

`UIWidget`<sup>API</sup>

`UIFormHeader`<sup>API</sup>

`UIFormBottomBar`<sup>API</sup>

## 89.4 Custom widgets

Custom widgets are included in the Java API to allow the development of user interface components for fields or groups of fields. A custom widget (`UICustomWidget`) allows, for a given node, to control the area where the input or display component is located. This allows having an input and display component that is fully customizable in HTML. The standard components (`UIWidgets`) are available and can be used. The custom widget can implement several display aspects: input component in the form, display in the form, display in a table cell. If a custom widget writes its own HTML components, it has the possibility to save the value in the database when submitting the form.

See also `UICustomWidget`<sup>API</sup>

## 89.5 Table filter

A table filter allows, for a given table, to create a criteria input form in order to apply a filter to the table view. The `UITableFilter` API is used to implement a table filter with a custom UI. It provides methods to create a UI that automatically adapts to the underlying data format (for example, by displaying a combo box when applicable).

See also

`UITableFilter`<sup>API</sup>

[Properties of data model elements](#) [p 51]

`UILabelRendererForHierarchy`<sup>API</sup>

## 89.6 Record label

EBX® uses a label to display a reference to a given record (for example a foreign key). Labels are also used in the title of a record form and in hierarchical views. This label can be customized in the model using expressions. It is also possible to customize labels using the `UILabelRenderer` API.

See also `UILabelRenderer`<sup>API</sup>

## 89.7 CSS and JavaScript

It is possible to integrate CSS and JavaScript files in each EBX® page by declaring them in the registration module. The inclusion of JavaScript files can be subject to conditions through development depending on the context.

### See also

[\*Module registration\*](#) [p 462]

[\*Development recommendations\*](#) [p 593]

*UIDependencyRegisterer*<sup>API</sup>

## CHAPTER 90

---

# Overview

A user service is an extension to TIBCO EBX® that provides a graphical user interface (GUI) allowing users to access specific or advanced functionalities.

An API is available allowing the development of powerful custom user services using the same visual components and data validation mechanisms as standard EBX® user interfaces.

This chapter contains the following topics:

1. [Nature](#)
2. [Declaration](#)
3. [Display](#)
4. [Legacy user services](#)

## 90.1 Nature

User services exist in different types called *natures*. The nature defines the minimal elements (dataspace, dataset, table, record...) that need to be selected to execute the service. The following table lists the available natures.

| Nature            | Description                                                                                                                                                                                                            |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Dataspace         | The nature of a user service that can be launched from the actions menu of a dataspace (branch or snapshot) or from any context where the current selection implies selecting a dataspace.                             |
| Dataset           | The nature of a user service that can be launched from the actions menu of a dataset or from any context where the current selection implies selecting a dataset.                                                      |
| TableView         | The nature of a user service that can be launched from the toolbar of a table, regardless of the selected view, or from any context where the current selection implies selecting a table.                             |
| Record            | The nature of a user service that can be launched from the toolbar of a record form or from any context where the current selection implies selecting a <i>single</i> record.                                          |
| Hierarchy         | The nature of a user service that can be launched from the toolbar of a table when a hierarchy view is selected.                                                                                                       |
| HierarchyNode     | The nature of a user service that can be launched from the menu of a table hierarchy view node. Currently, only <i>record</i> hierarchy nodes are supported.                                                           |
| Association       | The nature of a user service that can be launched from the target table view of an association or for any context where the current selection implies selecting the target table view of an association.               |
| AssociationRecord | The nature of a user service that can be launched from the form of a target record of an association node or from any context where the current selection implies selecting a <i>single</i> association target record. |

## 90.2 Declaration

A user service can be declared at two levels:

- Module,
- Data model.

A service declared by a data model can only be launched when the current selection includes a dataset of this model. The user service cannot be of the Dataspace nature.

A service declared by a module may be launched for any dataspace or dataset.

The declaration can add restrictions on selections that are valid for the user service.







## CHAPTER 91

---

# Quick start

This chapter contains the following topics:


1. [Main classes](#)
2. [Hello world](#)

## 91.1 Main classes

The minimum requirement is to implement two classes, one for the service declaration and one for the implementation itself.

## 91.2 Hello world

The sample is a dataset user service that simply displays a "hello" message, it can be launched from the action menu of a dataset:



The service implementation class must implement the interface `UserService<DatasetEntitySelection>`:

```
/**
 * This service displays hello world!
 */
public class HelloWorldService implements UserService<DatasetEntitySelection>
{
    public HelloWorldService()
    {
    }

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<DatasetEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        // Set bottom bar
        UIButtonSpecNavigation closeButton = aConfigurator.newCloseButton();
        closeButton.setDefaultButton(true);
        aConfigurator.setLeftButtons(closeButton);

        // Set content callback
        aConfigurator.setContent(this::writeHelloWorld);
    }

    private void writeHelloWorld(
        UserServicePaneContext aPaneContext,
```

```

    UserServicePaneWriter aWriter)
{
    // Display Hello World!

    aWriter.add("<div ");
    aWriter.addSafeAttribute("class", UICSSClasses.CONTAINER_WITH_TEXT_PADDING);
    aWriter.add(">");
    aWriter.add("Hello World!");
    aWriter.add("</div>");
}

@Override
public void setupObjectContext(
    UserServiceSetupObjectContext<DatasetEntitySelection> aContext,
    UserServiceObjectContextBuilder aBuilder)
{
    // No context yet.
}

@Override
public void validate(UserServiceValidateContext<DatasetEntitySelection> aContext)
{
    // No custom validation is necessary.
}

@Override
public UserServiceEventOutcome processEventOutcome(
    UserServiceProcessEventOutcomeContext<DatasetEntitySelection> aContext,
    UserServiceEventOutcome anEventOutcome)
{
    // By default do not modify the outcome.
    return anEventOutcome;
}
}

```

The declaration class must implement the interface `UserServiceDeclaration.OnDataset`:

```

/**
 * Declaration for service hello world!
 */
public class HelloWorldServiceDeclaration implements UserServiceDeclaration.OnDataset
{
    // The service key identifies the user service.
    private static final ServiceKey serviceKey = ServiceKey.forName("HelloWorld");

    public HelloWorldServiceDeclaration()
    {
    }

    @Override
    public ServiceKey getServiceKey()
    {
        return serviceKey;
    }

    @Override
    public UserService<DatasetEntitySelection> createUserService()
    {
        // Creates an instance of the user service.
        return new HelloWorldService();
    }

    @Override
    public void defineActivation(ActivationContextOnDataset aContext)
    {
        // The service is activated for all datasets instantiated with
        // the associated data model (see next example).
    }

    @Override
    public void defineProperties(UserServicePropertiesDefinitionContext aContext)
    {
        // This label is displayed in menus that can execute the user service.
        aContext.setLabel("Hello World Service");
    }

    @Override
    public void declareWebComponent(WebComponentDeclarationContext aContext)
    {
    }
}

```

In this sample, the user service is registered by a data model. The data model needs to define a schema extension that implements the following code:

```
public class CustomSchemaExtensions implements SchemaExtensions
{
    @Override
    public void defineExtensions(SchemaExtensionsContext aContext)
    {
        // Register the service.
        aContext.registerUserService(new HelloWorldServiceDeclaration());
    }
}
```

For details on the declaration of schema extensions, see `SchemaExtensionsAPI`.

## CHAPTER 92

# Implementing a user service

This chapter contains the following topics:

1. [Implementation interface](#)
2. [Life cycle and threading model](#)
3. [Object Context](#)
4. [Display setup](#)
5. [Database updates](#)
6. [Ajax](#)
7. [REST data services](#)
8. [File upload](#)
9. [File download](#)
10. [User service without display](#)

## 92.1 Implementation interface

The following table lists, per nature, the interface to implement:

| Nature            | Interface                                     |
|-------------------|-----------------------------------------------|
| Dataspace         | UserService<DataspaceEntitySelection>         |
| Dataset           | UserService<DatasetEntitySelection>           |
| TableView         | UserService<TableViewEntitySelection>         |
| Record            | UserService<RecordEntitySelection>            |
| Hierarchy         | UserService<HierarchyEntitySelection>         |
| HierarchyNode     | UserService<HierarchyNodeEntitySelection>     |
| Association       | UserService<AssociationEntitySelection>       |
| AssociationRecord | UserService<AssociationRecordEntitySelection> |

## 92.2 Life cycle and threading model

The user service implementation class is:

- Instantiated at the first HTTP request by a call to its declaration **createUserService()** `UserServiceDeclaration.createUserServiceAPI` method.
- Discarded when the current page goes out of scope or when the session times out.

Access to this class is synchronized by TIBCO EBX® to make sure that only one HTTP request is processed at a time. Therefore, the class does not need to be thread-safe.

The user service may have attributes. The state of these attributes will be preserved between HTTP requests. However, developers must be aware that these attributes should have moderate use of resources, such as memory, not to overload the EBX® server.

## 92.3 Object Context

The object context is a container for objects managed by the user service. This context is initialized and modified by the user service's implementation of the method `UserService.setupObjectContextAPI`.

An object of the object context is identified by an object key:

```
ObjectKey customerKey = ObjectKey.forName("customer");
```

An object can be:

- A record,
- A dataset,
- A new record not yet persisted,
- A dynamic object.

The object context is maintained between HTTP requests and usually only needs to be set up upon the first request.

Once persisted, a *new record* object is automatically changed to a *plain record* object.

As with **adaptations** `AdaptationAPI`, **path** `PathAPI` expressions are used to reference a sub-element of an object.

In the following sample, a pane writer adds a form input mapped to the attribute of an object:

```
// Add an input field for customer's last name.
aWriter.setCurrentObject(customerKey);
aWriter.addFormRow(Path.parse("lastName"));
```

In the following sample, an event callback gets the value of the attribute of an object:

```
// Get value of customer's last name.
ValueContext customerValueContext = aValueContext.getValueContext(customerKey);
String lastName = customerValueContext.getValue(Path.parse("lastName"));
```

A *dynamic object* is an object whose schema is defined by the user service itself. An API is provided to define the schema programmatically. This API allows defining only instance elements (instance nodes). Defining tables is not supported. It supports most other features available with standard EBX® data models, such as types, labels, custom widgets, enumerations and constraints, including programmatic ones.

The following sample defines two objects having the same schema:

```
public class SampleService implements UserService<TableViewEntitySelection>
{
    // Define an object key per object:
    private final static ObjectKey _PersonObjectKey = ObjectKey.forName("person");
    private final static ObjectKey _PartnerObjectKey = ObjectKey.forName("partner");

    // Define a path for each property:
    private final static Path _FirstName = Path.parse("firstName");
    private final static Path _LastName = Path.parse("lastName");
    private final static Path _BirthDate = Path.parse("birthDate");

    ...

    // Define and register objects:
    @Override
    public void setupObjectContext(
        UserServiceSetupObjectContext<DataspaceEntitySelection> aContext,
        UserServiceObjectContextBuilder aBuilder)
    {
        if (aContext.isInitialDisplay())
        {
            BeanDefinition def = aBuilder.createBeanDefinition();

            BeanElement firstName = def.createElement(_FirstName, SchemaTypeName.XS_STRING);
            firstName.setLabel("First name");
            firstName.setDescription("This is the given name");
            firstName.setMinOccurs(1);

            BeanElement lastName = def.createElement(_LastName, SchemaTypeName.XS_STRING);
            lastName.setLabel("Last name");
            lastName.setDescription("This is the family name");
            lastName.setMinOccurs(1);

            BeanElement birthDate = def.createElement(_BirthDate, SchemaTypeName.XS_DATE);
            birthDate.setLabel("Birth date");
            birthDate.addFacetMax(new Date(), false);

            aBuilder.registerBean(_PersonObjectKey, def);
            aBuilder.registerBean(_PartnerObjectKey, def);
        }

        ...
    }
}
```

## 92.4 Display setup

The display is set up by the user service's implementation of the method `UserService.setupDisplayAPI`. This method is called at each request and can set the following:

- The title (the default is the label specified by the user service declaration),
- The contextual help URL,
- The breadcrumbs,
- The toolbar,
- The bottom buttons.

If necessary, the header and the bottom buttons can be hidden.

The display setup is not persisted and, at each HTTP request, is reset to default before calling the method `UserService.setupDisplayAPI`.

### Bottom buttons

Buttons may be of two types: *action* and *submit*.

An *action* button triggers an *action* event without submitting the form. By default, the user needs to acknowledge that, by leaving the page, the last changes will be lost. This behavior can be customized.

A *submit* button triggers a *submit* event that always submits the form.

More information on events can be found in the following sections.


### Content callback

This callback usually implements the interface `UserServicePaneAPI` to render a plain EBX® form. The callback can also be an instance of `UserServiceRootTabbedPaneAPI` to render an EBX® form with tabs. For specific cases, the callback can implement `UserServiceRawPaneAPI`. This interface has restrictions but is useful when one wants to implement an HTML form that is not managed by EBX®.


### Toolbars

Toolbars are optional and come in two flavors.

The *form* style:



The *table* view style:



The style is automatically selected: toolbars defined for a *record* are of the form style and toolbars defined for a *table* are of the table view style.

### Samples

The following sample implements a button that closes the current user service and redirects the user back to the current selection, only if saving the data was successful:

```
public class SampleService implements UserService<...>
{
    private final static ObjectKey _RecordObjectKey = ObjectKey.forName("record");

    ...

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<RecordEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        ...
        // Define a "save and close" button with callback onSave().
        aConfigurator.setLeftButtons(aConfigurator.newSaveCloseButton(this::onSave));
    }

    private UserServiceEventOutcome onSave(UserServiceEventContext anEventContext)
```



```

{
    ProcedureResult result = anEventContext.save(_RecordObjectKey);
    if (result.hasFailed())
    {
        // Save has failed. Redisplay the user message.
        return null;
    }

    // Save has succeeded. Close the service.
    return UserServiceNext.nextClose();
}
}

```

The following sample is compatible with the Java 6 syntax. Only differences with the previous code are shown:

```

public class SampleService implements UserService<...>
{
    ...

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<RecordEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        ...
        // Define a "save and close" button with callback onSave().
        aConfigurator.setLeftButtons(aConfigurator.newSaveCloseButton(new UserServiceEvent() {
            @Override
            public UserServiceEventOutcome processEvent(UserServiceEventContext anEventContext)
            {
                return onSave(anEventContext);
            }
        }));
    }
}

```

The following sample implements a URL that closes the service and redirects the current user to another user service:

```

public class SampleService implements UserService<...>
{
    ...
    private void writePane(UserServicePaneContext aPaneContext, UserServicePaneWriter aWriter)
    {
        // Displays an URL that redirect current user.
        String url = aWriter.getURLForAction(this::goElsewhere);
        aWriter.add("<a ");
        aWriter.addSafeAttribute("href", url);
        aWriter.add(">Go elsewhere</a");
    }

    private UserServiceEventOutcome goElsewhere(UserServiceEventContext anEventContext)
    {
        // Redirects current user to another user service.
        ServiceKey serviceKey = ServiceKey.forModuleServiceName("CustomerModule", "CustomService");
        return UserServiceNext.nextService(serviceKey);
    }
}

```

The following code is an implementation of the method `UserService.processEventOutcome`<sup>APT</sup>, sufficient for simple user services:

```

public class HelloWorldService implements UserService<...>
{
    @Override
    public UserServiceEventOutcome processEventOutcome(
        UserServiceProcessEventOutcomeContext<DatasetEntitySelection> aContext,
        UserServiceEventOutcome anEventOutcome)
    {
        // By default do not modify the outcome.
        return anEventOutcome;
    }
}

```

The following sample is a more complex "wizard" service that includes three steps, each having its own `UserService.setupDisplay`<sup>APT</sup> method:

```

// Custom outcome values.

```

```

public enum CustomOutcome implements UserServiceEventOutcome {
    displayStep1, displayStep2, displayStep3
};

// All steps of the wizard service implement this interface.
public interface WizardStep
{
    public void setupDisplay(
        UserServiceSetupDisplayContext<DataspaceEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator);
}

// The user service implementation.
public class WizardService implements UserService<...>
{
    // Attribute for current step.
    private WizardStep step = new WizardStep1();

    ...

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<DataspaceEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        ...

        // Display current step.
        this.step.setupDisplay(aContext, aConfigurator);
    }

    @Override
    public UserServiceEventOutcome processEventOutcome(
        UserServiceProcessEventOutcomeContext<DataspaceEntitySelection> aContext,
        UserServiceEventOutcome anEventOutcome)
    {
        // Custom outcome value processing.

        if (anEventOutcome instanceof CustomOutcome)
        {
            CustomOutcome action = (CustomOutcome) anEventOutcome;
            switch (action)
            {
                case displayStep1:
                    this.step = new WizardStep1();
                    break;

                case displayStep2:
                    this.step = new WizardStep2();
                    break;

                case displayStep3:
                    this.step = new WizardStep3();
                    break;
            }

            // Redisplay the user service.
            return null;
        }

        // Let EBX® process the event outcome.
        return anEventOutcome;
    }
}

```

## 92.5 Database updates

An event callback may update the database.

The following sample saves two objects using a single transaction:

```

public class MultipleObjectsSampleService implements UserService<...>
{
    // This service defines a two objects having same schema.
    private final static ObjectKey _Person1_ObjectKey = ObjectKey.forName("person1");
    private final static ObjectKey _Person2_ObjectKey = ObjectKey.forName("person2");

    ...

    // Save button callback.

```

```
private UserServiceEventOutcome onSave(UserServiceEventContext aContext)
{
    ProcedureResult result = aContext.save(_Person1_ObjectKey, _Person2_ObjectKey);
    if (result.hasFailed())
    {
        //Save failed. Redisplay the service.
        //The user interface will automatically report error messages.
        return null;
    }

    // Save succeeded. Close the service.
    return UserServiceNext.nextClose();
}
}
```

The following sample updates the database using a **procedure** Procedure<sup>API</sup>:

```
import com.orchestranetworks.service.*;
import com.orchestranetworks.userservice.*;

public class MultipleObjectsSampleService implements UserService<...>
{
    ...

    // Event callback.
    private UserServiceEventOutcome onUpdateSomething(UserServiceEventContext aContext)
    {
        Procedure procedure = new Procedure()
        {
            public void execute(ProcedureContext aContext) throws Exception
            {
                // Code that updates database should be here.
                ...
            }
        };

        UserServiceTransaction transaction = aContext.createTransaction();
        transaction.add(procedure);

        ProcedureResult result = transaction.execute();
        if (result.hasFailed())
        {
            aContext.addError("Procedure failed");
        }
        else
        {
            aContext.addInfo("Procedure succeeded");
        }

        return null;
    }
}
```

## 92.6 Ajax

A user service can implement Ajax callbacks. An Ajax callback must implement the interface UserServiceAjaxRequest<sup>API</sup>.

The client calls an Ajax callback using the URL generated by: UserServiceResourceLocator.  
getURLForAjaxRequest<sup>API</sup>.

To facilitate the use of Ajax components, EBX® provides the JavaScript prototype  
EBX\_AJAXResponseHandler for sending the request and handling the response. For more information  
on EBX\_AJAXResponseHandler see UserServiceAjaxRequest<sup>API</sup>.

The following sample implements an Ajax callback that returns partial HTML:

```
public class AjaxSampleService implements UserService<DataspaceEntitySelection>
{
    ...

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<DataspaceEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        aConfigurator.setLeftButtons(aConfigurator.newCloseButton());
        aConfigurator.setContent(this::writePane);
    }
}
```

```

}

/**
 * Displays an URL that will execute the callback
 * and display the returned partial HTML inside a <div> tag.
 */
private void writePane(UserServicePaneContext aPaneContext, UserServicePaneWriter aWriter)
{
    // Generate the URL of the Ajax callback.
    String url = aWriter.getURLForAjaxRequest(this::ajaxCallback);

    // The id of the <div> that will display the partial HTML returned by the Ajax callback.
    String divId = "sampleId";

    aWriter.add("<div ");
    aWriter.addSafeAttribute("class", UICSSClasses.CONTAINER_WITH_TEXT_PADDING);
    aWriter.add(">");

    // Display the URL that will execute the callback.
    aWriter.add("<a ");
    aWriter.addSafeAttribute("href", "javascript:sample_sendAjaxRequest('" + url + "', '"
        + divId + "')");
    aWriter.add(">");
    aWriter.add("Click to call a user service Ajax callback");
    aWriter.add("</a>");

    // Output the <div> tag that will display the partial HTML returned by the callback.
    aWriter.add("<div ");
    aWriter.addSafeAttribute("id", divId);
    aWriter.add("></div>");

    aWriter.add("</div>");

    // JavaScript method that will send the Java request.
    aWriter.addJS_cr();
    aWriter.addJS_cr("function sample_sendAjaxRequest(url, targetDivId) {");
    aWriter.addJS_cr("    var ajaxHandler = new EBX_AJAXResponseHandler();");

    aWriter.addJS_cr("    ajaxHandler.handleAjaxResponseSuccess = function(responseContent) {");
    aWriter.addJS_cr("        var element = document.getElementById(targetDivId);");
    aWriter.addJS_cr("        element.innerHTML = responseContent;");
    aWriter.addJS_cr("    }");

    aWriter.addJS_cr("    ajaxHandler.handleAjaxResponseFailed = function(responseContent) {");
    aWriter.addJS_cr("        var element = document.getElementById(targetDivId);");
    aWriter.addJS_cr("        element.innerHTML = \"<span class='\" + UICSSClasses.TEXT.ERROR
        + \">Ajax call failed</span>\";");
    aWriter.addJS_cr("    }");

    aWriter.addJS_cr("    ajaxHandler.sendRequest(url);");
    aWriter.addJS_cr("}");
}

/**
 * The Ajax callback that returns partial HTML.
 */
private void ajaxCallback(
    UserServiceAjaxContext anAjaxContext,
    UserServiceAjaxResponse anAjaxResponse)
{
    UserServiceWriter writer = anAjaxResponse.getWriter();
    writer.add("<p style='\"color:green\">Ajax callback succeeded!</p>");
    writer.add("<p>Current data and time is: ");

    DateFormat format = DateFormat.getDateTimeInstance(
        DateFormat.FULL,
        DateFormat.FULL,
        Locale.US);
    writer.addSafeInnerHTML(format.format(new Date()));

    writer.add("</p>");
}
}

```

## 92.7 REST data services

A user service can access REST data services through HTTP requests.

The client should use the URL generated by: `UIResourceLocator.getURLForRestAPI`. This URL includes required information for the user authentication.

For more information on REST data services see the [Built-in RESTful services](#) [p 659].

The following sample implements a REST data service call whose response is printed in a textarea:

```
public class RestCallSampleService implements UserService<DataspaceEntitySelection>
{
    ...

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<DataspaceEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        aConfigurator.setLeftButtons(aConfigurator.newCloseButton());
        aConfigurator.setContent(this::writePane);
    }

    private void writePane(UserServicePaneContext aPaneContext, UserServicePaneWriter aWriter)
    {
        // Generates the URL for REST data service call without additional parameters
        final String url = aWriter.getURLForRest("/ebx-dataseservices/rest/{specificPath}", null);

        final String resultAreaId = "restResult";

        // Displays a link for REST data service call
        aWriter.add("<div ");
        aWriter.addSafeAttribute("class", UICSSClasses.CONTAINER_WITH_TEXT_PADDING);
        aWriter.add(">");
        aWriter.add("<p>This link will display the response after making a REST call</p>");
        aWriter.add("<a ");
        aWriter.addSafeAttribute("href",
            "javascript:sendRestRequest('" + url + "', '" + resultAreaId + "')");
        aWriter.add(">");
        aWriter.add("Make the call.");
        aWriter.add("</a>");
        aWriter.add("<textarea ");
        aWriter.addSafeAttribute("id", resultAreaId);
        aWriter.add(" readonly=\"readonly\" style=\"width: 100%;\" ></textarea>");
        aWriter.add("</div>");

        // JavaScript method that will send the HTTP REST request
        aWriter.addJS_cr("function sendRestRequest(url, targetId) {");
        aWriter.addJS_cr("    var xhttp = new XMLHttpRequest();");
        aWriter.addJS_cr("    xhttp.open('GET', url, true);");
        aWriter.addJS_cr("    xhttp.setRequestHeader('Content-type', 'application/json');");
        aWriter.addJS_cr("    xhttp.send();");
        aWriter.addJS_cr("    var element = document.getElementById(targetId);");
        aWriter.addJS_cr("    xhttp.onreadystatechange = function() {");
        aWriter.addJS_cr("        if (xhttp.readyState == 4)");
        aWriter.addJS_cr("            element.innerHTML = xhttp.responseText;");
        aWriter.addJS_cr("    }");
        aWriter.addJS_cr("}");
    }
}
```

## 92.8 File upload

A user service can display forms with file input fields.

The following sample displays a form with two input fields, a title and a file:

```
public class FileUploadService implements UserService<...>
{
    // This service defines a single object named "file".
    private final static ObjectKey _File_ObjectKey = ObjectKey.forName("file");

    // Paths for the "file" object.
    public final static Path _Title = Path.parse("title");
    public final static Path _File = Path.parse("file");

    ...

    @Override
    public void setupObjectContext(
        UserServiceSetupObjectContext<DataspaceEntitySelection> aContext,
        UserServiceObjectContextBuilder aBuilder)
    {
        if (aContext.isInitialDisplay())
        {
            // Create a definition for the "model" object.
        }
    }
}
```

```

BeanDefinition def = aBuilder.createBeanDefinition();
aBuilder.registerBean(_File_ObjectKey, def);

BeanElement element;

element = def.createElement(_Title, SchemaTypeName.XS_STRING);
element.setLabel("Title");
element.setMinOccurs(1);

// Type for a file must be BeanDefinition.OSD_FILE_UPLOAD.
element = def.createElement(_File, BeanDefinition.OSD_FILE_UPLOAD);
element.setLabel("File");
element.setMinOccurs(1);
}
}

@Override
public void setupDisplay(
    UserServiceSetupDisplayContext<DataspaceEntitySelection> aContext,
    UserServiceDisplayConfigurator aConfigurator)
{
    aConfigurator.setTitle("File upload service");
    aConfigurator.setLeftButtons(aConfigurator.newSubmitButton("Upload", this::onUpload),
    aConfigurator.newCancelButton());

    // IMPORTANT: Following method must be called to enable file upload.
    // This will set form encryption type to "multipart/form-data".
    aConfigurator.setFileUploadEnabled(true);

    aConfigurator.setContent(this::writePane);
}

private void writePane(UserServicePaneContext aContext, UserServicePaneWriter aWriter)
{
    final UIWidgetFileUploadFactory fileUploadFactory = new UIWidgetFileUploadFactory();

    aWriter.setCurrentObject(_File_ObjectKey);

    aWriter.startTableFormRow();

    // Title input.
    aWriter.addFormRow(_Title);

    // File upload input.
    UIWidgetFileUpload widget = aWriter.newCustomWidget(_File, fileUploadFactory);
    // Default filter for file names.
    widget.setAccept(".txt");
    aWriter.addFormRow(widget);

    aWriter.endTableFormRow();
}

private UserServiceEventOutcome onUpload(UserServiceEventContext anEventContext)
{
    ValueContextForInputValidation valueContext = anEventContext.getValueContext(_File_ObjectKey);

    String title = (String) valueContext.getValue(_Title);
    UploadedFile file = (UploadedFile) valueContext.getValue(_File);

    InputStream in;
    try
    {
        in = file.getInputStream();
    }
    catch (IOException e)
    {
        // Should not happen.
        anEventContext.addError("Cannot read file.");
        return null;
    }

    // Do something with title and the input stream.
    return UserServiceNext.nextClose();
}
}

```

For more information, see `UIWidgetFileUploadAPT`.

## 92.9 File download

A user service can display URLs or buttons to download files. The actual downloading of a file is under the control of the user service.

The following sample displays a URL to download a file:

```
public class FileDownloadService implements UserService<DataspaceEntitySelection>
{
    ...

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<DataspaceEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        aConfigurator.setLeftButtons(aConfigurator.newCloseButton());
        aConfigurator.setContent(this::writePane);
    }

    private void writePane(UserServicePaneContext aContext, UserServicePaneWriter aWriter)
    {
        aWriter.add("<div ");
        aWriter.addSafeAttribute("class", UICSSClasses.CONTAINER_WITH_TEXT_PADDING);
        aWriter.add(">");

        // Generate and display the URL for the download.
        String downloadURL = aWriter.getURLForGetRequest(this::processDownloadRequest);

        aWriter.add("<a ");
        aWriter.addSafeAttribute("href", downloadURL);
        aWriter.add(">Click here to download a sample file</a>");

        aWriter.add("</div>");
    }

    private void processDownloadRequest(
        UserServiceGetContext aContext,
        UserServiceGetResponse aResponse)
    {
        // The file is plain text.
        aResponse.setContentType("text/plain;charset=UTF-8");
        // Remove the following statement to display the file directly in the browser.
        aResponse.setHeader("Content-Disposition", "attachment; filename=\"sample.txt\"");

        // Write a text file using UTF-8 encoding.
        PrintWriter out;
        try
        {
            out = new PrintWriter(new OutputStreamWriter(aResponse.getOutputStream(), "UTF-8"));
        }
        catch (IOException ex)
        {
            throw new RuntimeException(ex);
        }

        DateFormat format = DateFormat.getDateInstance(
            DateFormat.FULL,
            DateFormat.MEDIUM,
            Locale.US);
        Date now = new Date();

        out.println("Hello !");
        out.println("This is a sample text file downloaded on " + format.format(now)
            + ", from EBX®.");

        out.close();
    }
}
```

## 92.10 User service without display

A user service may be designed to execute a task without display and return to the previous screen or redirect the user to another screen.

This type of service must implement the interface **UserServiceExtended** `UserServiceExtendedAPI` and method `UserServiceExtended.initializeAPI`.

The following sample deletes selected records in the current table view:

```
public class DeleteRecordsService implements UserServiceExtended<TableViewEntitySelection>
{
    ...

    @Override
    public UserServiceEventOutcome initialize(
        UserServiceInitializeContext<TableViewEntitySelection> aContext)
    {
        final List<AdaptationName> records = new ArrayList<>();

        // Deletes all selected rows in a single transaction.
        RequestResult requestResult = aContext.getEntitySelection().getSelectedRecords().execute();
        try
        {
            for (Adaptation record = requestResult.nextAdaptation(); record != null; record =
requestResult.nextAdaptation())
            {
                records.add(record.getAdaptationName());
            }
        }
        finally
        {
            requestResult.close();
        }

        Procedure deleteProcedure = new Procedure()
        {
            @Override
            public void execute(ProcedureContext aContext) throws Exception
            {
                for (AdaptationName record : records)
                {
                    aContext.doDelete(record, false);
                }
            }
        };

        UserServiceTransaction transaction = aContext.createTransaction();
        transaction.add(deleteProcedure);

        // Adds an information messages for current user.
        ProcedureResult procedureResult = transaction.execute(true);
        if (!procedureResult.hasFailed())
        {
            if (records.size() <= 1)
            {
                aContext.addInfo(records.size() + " record was deleted.");
            }
            else
            {
                aContext.addInfo(records.size() + " records were deleted.");
            }
        }

        // Do not display the user service and return to current view.
        return UserServiceNext.nextClose();
    }

    @Override
    public void setupObjectContext(
        UserServiceSetupObjectContext<TableViewEntitySelection> aContext,
        UserServiceObjectContextBuilder aBuilder)
    {
        //Do nothing.
    }

    @Override
    public void setupDisplay(
        UserServiceSetupDisplayContext<TableViewEntitySelection> aContext,
        UserServiceDisplayConfigurator aConfigurator)
    {
        //Do nothing.
    }

    @Override
    public void validate(UserServiceValidateContext<TableViewEntitySelection> aContext)
    {
        //Do nothing.
    }
}
```



```
}  
  
@Override  
public UserServiceEventOutcome processEventOutcome(  
    UserServiceProcessEventOutcomeContext<TableViewEntitySelection> aContext,  
    UserServiceEventOutcome anEventOutcome)  
{  
    return anEventOutcome;  
}  
}
```

### ***Known limitation***

If such service is called in the context of a Web component, an association, a perspective action or a hierarchy node, The service will be launched, initialized and closed, but the service's target entity will still be displayed.



## CHAPTER 93

# Declaring a user service

This chapter contains the following topics:

1. [Declaration interface](#)
2. [Life cycle and threading model](#)
3. [Registration](#)
4. [Service properties](#)
5. [Service activation scope](#)
6. [Web component declaration](#)
7. [User service groups](#)

## 93.1 Declaration interface

The following table lists, per nature, the interface that the declaration class of a user service must implement:

| Nature            | Declaration Interface                      |
|-------------------|--------------------------------------------|
| Dataspace         | UserServiceDeclaration.OnDataspace         |
| Dataset           | UserServiceDeclaration.OnDataset           |
| TableView         | UserServiceDeclaration.OnTableView         |
| Record            | UserServiceDeclaration.OnRecord            |
| Hierarchy         | UserServiceDeclaration.OnHierarchy         |
| HierarchyNode     | UserServiceDeclaration.OnHierarchyNode     |
| Association       | UserServiceDeclaration.OnAssociation       |
| AssociationRecord | UserServiceDeclaration.OnAssociationRecord |

## 93.2 Life cycle and threading model

The user service declaration class is instantiated at the TIBCO EBX® startup and must be coded to be thread-safe. This is usually not an issue as most implementations should be immutable classes.

## 93.3 Registration

A user service declaration must be registered by a module or a data model.

Registration by a module is achieved by the module registration servlet by a code similar to:

```
public class CustomRegistrationServlet extends ModuleRegistrationServlet
{
    @Override
    public void handleServiceRegistration(ModuleServiceRegistrationContext aContext)
    {
        // Register custom user service declaration.
        aContext.registerUserService(new CustomServiceDeclaration());
    }
}
```

For more information on the module registration servlet, see [module registration](#) [p 462] and `ModuleRegistrationServlet`<sup>API</sup>.

Registration by a data model is achieved by a code similar to:

```
public class CustomSchemaExtensions implements SchemaExtensions
{
    @Override
    public void defineExtensions(SchemaExtensionsContext aContext)
    {
        // Register custom user service declaration.
        aContext.registerUserService(new CustomServiceDeclaration());
    }
}
```

For more information on data model extensions, see `SchemaExtensions`<sup>API</sup>.

## 93.4 Service properties

The properties of a user service include its *label*, *description*, *confirmation message* and the *group* that owns the service. All are optional but it is a good practice to at least define the label.

For more information, see `UserServiceDeclaration.defineProperties`<sup>API</sup>.

## 93.5 Service activation scope

The activation scope defines on which selection the service is available.

Example of a service activation definition:

```
public class CustomServiceDeclaration implements UserServiceDeclaration.OnTableView
{
    ...

    @Override
    public void defineActivation(ActivationContextOnTableView aContext)
    {
        // activates the service in all dataspace except the "Reference" branch.
        aContext.includeAllDataspaces(DataspaceType.BRANCH);
        aContext.excludeDataspacesMatching(Repository.REFERENCE, DataspaceChildrenPolicy.NONE);

        // activates the service only on tables "table01" and "table03".
        aContext.includeSchemaNodesMatching(
            CustomDataModelPath._Root_Table01.getPathInSchema(),
            CustomDataModelPath._Root_Table03.getPathInSchema());

        // service will be enabled only when at least one record is selected.
    }
}
```

```

aContext.forbidEmptyRecordSelection();

// service will not be displayed in hierarchical views (neither in the
// top toolbar, nor in the hierarchy nodes' menu).
aContext.setDisplayForLocations(
    ActionDisplaySpec.HIDDEN,
    ToolbarLocation.HIERARCHICAL_VIEW_TOP,
    ToolbarLocation.HIERARCHICAL_VIEW_NODE);

// service will be considered as disabled if not explicitly enabled
// via the UI.
aContext.setDefaultPermission(UserServicePermission.getDisabled());
}
}

```

For more information about declaring the activation scope, see `UserServiceDeclaration.defineActivationAPI`.

For more information about the resolution of the user service availability, see [Resolving permissions on services](#) [p 285].

## 93.6 Web component declaration

### *Parameters declaration and availability in workflows and perspectives*

User services are automatically available as web components with a set of built-in parameters depending on the service's nature. To define custom parameters and/or set the service web component as available when configuring a workflow user task, a perspective menu action or a toolbar web component action, `UserServiceDeclaration.declareWebComponentAPI` must be used.

Example of a web component declaration:

```

public class CustomServiceDeclaration implements UserServiceDeclaration.OnDataset
{
    ...

    @Override
    public void declareWebComponent(WebComponentDeclarationContext aContext)
    {
        // makes this web component available when configuring a workflow user task.
        aContext.setAvailableAsWorkflowUserTask(true);

        // adds a custom input parameter.
        aContext.addInputParameter(
            "source",
            UserMessage.createInfo("Source"),
            UserMessage.createInfo("Source of the imported data.));

        // modifies the built-in "instance" parameter to be "input/output" instead of "input".
        aContext.getBuiltInParameterForOverride("instance").setOutput(true);
    }
}

```

See [Using TIBCO EBX® as a Web Component](#) [p 191] for more information.

### *User service extension*

It is possible to extend existing user services (built-in or custom) in order to add input/output parameters when using these services as web components.

In order to do so, a user service extension must first be registered by a module or a data model.

Registration by a module is achieved by the module registration servlet by code similar to:

```

public class CustomRegistrationServlet extends ModuleRegistrationServlet
{
    ...

    @Override
    public void handleServiceRegistration(ModuleServiceRegistrationContext aContext)
    {
        ...
    }
}

```

```
{
    // Register user service extension declaration.
    aContext.registerUserServiceExtension(new ServiceExtensionDeclaration());
}
}
```

For more information on the module registration servlet, see [module registration](#) [p 462] and `ModuleRegistrationServlet`<sup>API</sup>.

Registration by a data model is achieved by a code similar to:

```
public class CustomSchemaExtensions implements SchemaExtensions
{
    ...

    @Override
    public void defineExtensions(SchemaExtensionsContext aContext)
    {
        // Register user service extension declaration.
        aContext.registerUserServiceExtension(new ServiceExtensionDeclaration());
    }
}
```

For more information on the data model extension, see `SchemaExtensions`<sup>API</sup>.

## 93.7 User service groups

User service groups are used to organize the display of user services in menus and permission management screens.

The following types of service groups are available:

- [Built-in User Service Groups](#) [p 590] provided by EBX®,
- [Custom User Service Groups](#) [p 591] declared in a module.

The link between groups and services is made upon service declaration. See [Associating a service to a group](#) [p 591].

### ***Built-in User Service Groups***

Available built-in service groups:

| Service Group Key        | Description                                                                                                                                                                                                                                                                                                                      |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>@ebx-importExport</b> | Group containing all built-in import and export services provided by EBX®. In the default menus, these services are displayed in an "Import / Export" sub-menu.                                                                                                                                                                  |
| <b>@ebx-views</b>        | Group containing services to display in the 'Views' menu. Unlike other service groups, services associated with this group are not displayed in the default menus, but only in the 'Views' menu displayed in the non-customizable part of the table top toolbar. These services can still be added manually to a custom toolbar. |

## Declaring a User Service Group

User Service Groups must be declared while registering the module, using the method `ModuleServiceRegistrationContext.registerServiceGroup`<sup>API</sup>:

```
public class CustomRegistrationServlet extends ModuleRegistrationServlet
{
    ...

    @Override
    public void handleServiceRegistration(ModuleServiceRegistrationContext aContext)
    {
        // In CustomModuleConstants,
        // CUSTOM_SERVICE_GROUP_KEY = ServiceGroupKey.forServiceGroupInModule("customModule", "customGroup")

        // registers CUSTOM_SERVICE_GROUP_KEY service group
        aContext.registerServiceGroup(
            CustomModuleConstants.CUSTOM_SERVICE_GROUP_KEY,
            UserMessage.createInfo("Custom group"),
            UserMessage.createInfo("This group contains services related to..."));
    }
}
```

## Associating a service to a group

The association of a service with a group is made at its **declaration** `UserServiceDeclaration`<sup>API</sup>, using the method `UserServicePropertiesDefinitionContext.setGroup`<sup>API</sup>:

```
public class CustomServiceDeclaration implements UserServiceDeclaration.OnDataset
{
    ...

    @Override
    public void defineProperties(UserServicePropertiesDefinitionContext aContext)
    {
        // associates the current service to the CUSTOM_SERVICE_GROUP_KEY group
        aContext.setGroup(CustomModuleConstants.CUSTOM_SERVICE_GROUP_KEY);
    }
}
```

A service can be associated with either a built-in or a custom service group. In the latter case, this service will be displayed in this built-in group, just like other built-in services belonging to this group.





## CHAPTER 94

# Development recommendations

This chapter contains the following topics:

1. [HTML](#)
2. [CSS](#)
3. [JavaScript](#)

## 94.1 HTML

It is recommended to minimize the inclusion of specific HTML styles and tags to allow the default styles of TIBCO EBX® to apply to custom interfaces. The approach of the API is to automatically apply a standardized style to all elements on HTML pages, while simplifying the implementation process for the developer.

### *XHTML*

EBX® is a Rich Internet Application developed in XHTML 1.0 Transitional. It means that the structure of the HTML is strict XML file and that all tags must be closed, including "br" tags. This structure allows for greater control over CSS rules, with fewer differences in browser rendering.

### *iFrames*

Using iFrame is allowed in EBX®, especially in collaboration with a URL of a `UIHttpManagerComponentAPI`. For technical reasons, it is advised to set the `src` attribute of an iFrame using JavaScript only. In this way, the iFrame will be loaded once the page is fully rendered and when all the built-in HTML components are ready.

### Example

The following example, developed from any `UIComponentWriterAPI`, uses a `UIHttpManagerComponentAPI` to build the URL of an iFrame, and set it in the right way:

```
// using iFrame in the current page requires a sub session component
UIHttpManagerComponent managerComponent = writer.createWebComponentForSubSession();

// [...] managerComponent configuration

String iFrameURL = managerComponent.getURIWithParameters();

String iFrameId = "mySweetIFrame";

// place the iFrame in the page, with an empty src attribute
writer.add("<iframe id=\"\">".add(iFrameId).add("\n src=\"\" >").add("</iframe>");

// launch the iFrame from JavaScript
```

```
writer.addJS("document.getElementById(\"").addJS(iFrameId).addJS(\"").src = \"").addJS(iFrameURL).addJS(\"");");
```

## 94.2 CSS

### Public CSS classes

The constant catalog `UICSSClasses`<sup>API</sup> offers the main CSS classes used in the software to style the components. These CSS classes ensure a proper long-term integration into the software, because they follow the background colors, borders, customizable text in the administration; the floating margins and paddings fluctuate according to the variable density; to the style of the icons, etc.

See also `UICSSUtils`<sup>API</sup>

### Advanced CSS

EBX® allows to integrate to all its pages one or more external Cascading Style Sheet. These external CSS, considered as resources, need to be declared in the [Module registration](#) [p 462].

In order to ensure the proper functioning of your CSS rules and properties without altering the software, the following recommendations should be respected. Failure to respect these rules could lead to:

- Improper functioning of the software, both aesthetically and functionally: risk of losing the display of some of the data and some input components may stop working.
- Improper functioning of your CSS rules and properties, since the native CSS rules will impact the CSS implementation.

### Reserved prefixes for CSS identifiers and class names

The following prefixes should not be used to create CSS #ids and .classes.

|                   |                                                        |
|-------------------|--------------------------------------------------------|
| <b>ebx_</b>       | Internal built-in                                      |
| <b>yui</b>        | Yahoo User Interface global                            |
| <b>ygtv</b>       | Yahoo User Interface tree view                         |
| <b>layout-doc</b> | Yahoo User Interface layout                            |
| <b>cke_</b>       | CK editor (used by HTML editor widget)                 |
| <b>fa</b>         | Font Awesome (icons used by perspectives and toolbars) |

### CSS classes used internally by EBX®

The following CSS classes should never be included in a ruleset that has no contextual selector.

If you do not prefix your CSS selector using one of the CSS classes below, it will cause conflicts and corrupt the UI of EBX®.

|                        |                                  |
|------------------------|----------------------------------|
| <b>selected</b>        | YUI selected tree node           |
| <b>hd</b>              | YUI floating pane header         |
| <b>bd</b>              | YUI floating pane body           |
| <b>ft</b>              | YUI floating pane footer         |
| <b>container-close</b> | YUI inner popup close button     |
| <b>underlay</b>        | YUI inner popup shadow           |
| <b>hastitle</b>        | YUI menu group with title        |
| <b>topscrollbar</b>    | YUI menu top scroll zone         |
| <b>bottomscrollbar</b> | YUI menu bottom scroll zone      |
| <b>withtitle</b>       | YUI calendar                     |
| <b>link-close</b>      | YUI calendar close button        |
| <b>collapse</b>        | YUI layout closed pane indicator |
| <b>pull-right</b>      | Font Awesome parameter           |
| <b>pull-left</b>       | Font Awesome parameter           |

### Examples to avoid conflicts

#### *Don't*

```
.selected {
  background-color: red;
}
```

#### *Do*

```
#myCustomComponent li.selected {
  background-color: red;
}
```

## 94.3 JavaScript

### **Public JS functions**

The catalog of JavaScript functions `JavaScriptCatalogAPI` offers a list of functions to use directly (through copy-paste) in the JS files.

### **JavaScript call during page generation in Java**

When generating the HTML of a Java component, it is possible to add specific JavaScript code with the API `UIJavaScriptWriterAPI`.

This JavaScript is executed once the whole page is loaded. It is possible to instantly manage the HTML elements written with `UIBodyWriter.addAPI`. Setting on-load functions (such as `window.onload = myFunctionToCallOnLoad;`) is not supported because the execution context comes after the on-load event.

### **Advanced JavaScript**

EBX® allows to include one or more external JavaScript files. These external JavaScript files, considered as resources, need to be declared in the [Module registration](#) [p 462]. For performance reasons, it is recommended to include the JavaScript resource only when necessary (in a User service or a specific form, for example). The API `UIDependencyRegistererAPI` allows a developer to specify the conditions for which the JavaScript resources will be integrated into a given page according to its context.

In order to ensure the proper functioning of your JavaScript resources without altering the software, the following recommendations should be respected. Failure to respect them could lead to:

- Improper functioning of the software: if functions or global variables of the software were to be erased, some input or display components (including the whole screen) may stop working.
- Improper functioning of your JavaScript instructions, since global variables or function names could be erased.

### **Reserved JS prefixes**

The following prefixes are reserved and should not be used to create variables, functions, methods, classes, etc.

|              |                          |
|--------------|--------------------------|
| <b>ebx_</b>  | Internal built-in API    |
| <b>EBX_</b>  | Internal built-in API    |
| <b>YAHOO</b> | Yahoo User Interface API |

---

# SOAP data services

---

## CHAPTER 95

# Introduction

This chapter contains the following topics:

1. [Overview](#)
2. [Activation and configuration](#)
3. [Interactions](#)
4. [Security](#)
5. [Monitoring](#)
6. [SOAP and REST comparative](#)
7. [Limitations](#)

## 95.1 Overview

Data services allow external systems to interact with the data governed in the TIBCO EBX® repository using the SOAP/Web Services Description Language (WSDL) standards.

In order to invoke [SOAP operations](#) [p 617], for an integration use case, a [WSDL](#) [p 609] must be generated from a data model. It will be possible to perform operations such as:

- Selecting, inserting, updating, deleting, or counting records
- Selecting dataset values
- Getting the differences on a table between dataspace or snapshots, or between two datasets based on the same data model
- Getting the credentials of records

Other generic WSDLs can be generated and allow performing operations such as:

- Creating, merging, or closing a dataspace
- Creating or closing a snapshot
- Validating a dataset, dataspace, or a snapshot
- Starting, resuming or ending a data workflow

- Administrative operations to manage access to the UI or to system information

#### Note

See [SOAP and REST comparative](#) [p 606].

## 95.2 Activation and configuration

Data services are enabled by deploying the ebx-dataservices web application along with the other EBX® modules. See [Java EE deployment overview](#) [p 315] for more information.

In case of specific deployment, for example using reverse-proxy mode, see [URLs computing](#) [p 359] for more information.

The default method for accessing data services is over HTTP, although it is also possible to use JMS for the SOAP operations. See [JMS configuration](#) [p 358] and [Using JMS](#) [p 600] for more information.

## 95.3 Interactions

### *Input and output message encoding*

All input messages must be *exclusively* in UTF-8. All output messages are in UTF-8.

### *Tracking information*

Depending on the data services operation being called, it may be possible to specify session tracking information.

- Example for a SOAP operation, the request header contains:

```
<SOAP-ENV:Header>
  <!-- optional security header here -->
  <m:session xmlns:m="urn:ebx-schemas:dataservices_1.0">
    <trackingInformation>String</trackingInformation>
  </m:session>
</SOAP-ENV:Header>
```

For more information, see `Session.getTrackingInfoAPI` in the Java API.

### *Session parameters*

Depending on the data services operation being called, it is possible to specify session input parameters. They are defined in the request body.

Input parameters are available on custom Java components with a session object, such as: triggers, access rules, custom web services. They are also available on data workflow operations.

- Example for a SOAP operation, the optional request header contains:

```
<SOAP-ENV:Header>
  <!-- optional security header here -->
  <m:session xmlns:m="urn:ebx-schemas:dataservices_1.0">
    <!-- optional trackingInformation header here -->
    <inputParameters>
      <parameter>
        <name>String</name>
        <value>String</value>
      </parameter>
      <!-- for some other parameters, copy complex
           element 'parameter' -->
    </inputParameters>
  </m:session>
</SOAP-ENV:Header>
```

For more information, see `Session.getInputParameterValueAPI` in the Java API.

## Exception handling

In case of unexpected server error upon execution of:

- A SOAP operation, a SOAP exception response is returned to the caller via the `soap:Fault` element. For example:

```
<soapenv:Fault>
  <faultcode>soapenv:java.lang.IllegalArgumentException</faultcode>
  <faultstring />
  <faultactor>admin</faultactor>
  <detail>
    <m:StandardException xmlns:m="urn:ebx-schemas:dataservices_1.0">
      <code>java.lang.IllegalArgumentException</code>
      <label/>
      <description>java.lang.IllegalArgumentException:
        Parent home not found at
        com.orchestranetworks.XX.YY.ZZ.AA.BB(AA.java:44) at
        com.orchestranetworks.XX.YY.ZZ.CC.DD(CC.java:40) ...
      </description>
    </m:StandardException>
  </detail>
</soapenv:Fault>
```

## Using JMS

It is possible to access SOAP operations using JMS instead of HTTP. The JMS architecture relies on one JMS request queue (mandatory), on one JMS failure queue (optional), and on JMS response queues, see configuration [JMS](#) [p 358]. The mandatory queue is the input queue. Request messages must be put in the input queue, and response messages are put by EBX® in the `replyTo` queue of the JMS request. The optional queue is the failure queue which allows you to replay an input message if necessary. If the queue is set and activated in the configuration file and an exception occurs while handling a request message, this input message will be copied in the failure queue.

The relationship between a request and a response is made by copying the `messageId` message identifier field of the JMS request into the `correlId` correlation identifier field of the response.

JMS location points must be defined in the Lineage administration in order to specialize the generated WSDL. If no specific location point is given, the default value will be `jms:queue:jms/EBX_QueueIn`.

# 95.4 Security

## Authentication

Authentication is mandatory to access to data. Several authentication methods are available and described below. The descriptions are ordered by priority (EBX® applies the highest priority authentication method first).

- 'Basic Authentication Scheme' method is based on the HTTP-Header Authorization in base 64 encoding, as described in [RFC 2617 \(Basic Authentication Scheme\)](#).

```
If the user agent wishes to send the userid "Alibaba" and password "open sesame",
it will use the following header field:
> Authorization: Basic QWxpYmFiYTpvGvUHNlc2FtZQ==
```

- 'Standard Authentication Scheme' is based on the HTTP Request. User and password are extracted from request parameters. For more information on request parameters, see [Parameters](#) [p 612] section.



For more information on this authentication scheme, see `Directory.authenticateUserFromLoginPasswordAPI`.

- The 'SOAP Security Header Authentication Scheme' method is based on the [Web Services Security UsernameToken Profile 1.0](#) specification.

By default, the type `PasswordText` is supported. This is done with the following SOAP-Header defined in the WSDL:

```
<SOAP-ENV:Header>
  <wsse:Security xmlns:wsse="http://schemas.xmlsoap.org/ws/2002/04/secext">
    <wsse:UsernameToken>
      <wsse:Username>String</wsse:Username>
      <wsse:Password Type="wsse:PasswordText">String</wsse:Password>
    </wsse:UsernameToken>
  </wsse:Security>
</SOAP-ENV:Header>
```

#### Note

Only available for [SOAP operations](#) [p 617].

- 'Specific authentication Scheme' is based on the HTTP Request. An implementation of this method can, for example, extract a password-digest or a ticket from the HTTP request. See `Directory.authenticateUserFromHttpRequestAPI` for more information.
- The 'SOAP Specific Header Authentication Scheme'.

For more information, see [Overriding the SOAP security header](#) [p 601].

## Global permissions

Global access permissions can be independently defined for the SOAP and WSDL connector accesses. For more information see [Global permissions](#) [p 385].

## Overriding the SOAP security header

It is possible to override the default WSS header in order to define another security authentication mechanism. Such an override is taken into account for both HTTP and JMS. To define and override,

use the 'SOAP Header Security declaration' configuration settings under Administration > Lineage, which includes the following fields:

|                               |                                                                                                            |
|-------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Schema location</b>        | The URI of the Security XML Schema to import into the WSDL.                                                |
| <b>Target namespace</b>       | The target namespace of elements in the schema.                                                            |
| <b>Namespace prefix</b>       | The prefix for the target namespace.                                                                       |
| <b>Message name</b>           | The message name to use in the WSDL.                                                                       |
| <b>Root element name</b>      | The root element name of the security header. The name must be the same as the one declared in the schema. |
| <b>wSDL:part element name</b> | The name of the wSDL:part of the message.                                                                  |

The purpose of overriding the default security header is to change the declaration of the WSDL message matching the security header so that it contains the following:

```
<wSDL:definitions ... xmlns:MyPrefix="MyTargetNameSpace" ...
...
<xs:schema ...>
  <xs:import namespace="MyTargetNameSpace" schemaLocation="MySchemaURI"/>
  ...
</xs:schema>
...
<wSDL:message name="MySecurityMessage">
  <wSDL:part name="MyPartElementName" element="MyPrefix:MySecurityRootElement"/>
</wSDL:message>
...
<wSDL:operation name="...">
  <soap:operation soapAction="..." style="document"/>
  <wSDL:input>
    <soap:body use="literal"/>
    <soap:header message="impl:MySecurityMessage" part="MyPartElementName" use="literal"/>
  ...
</wSDL:operation>
</wSDL:definitions>
```

The corresponding XML Schema header declaration would be as follows:

```
<schema xmlns="http://www.w3.org/2001/XMLSchema" targetNamespace="MyNameSpace"
  xmlns:MyPrefix="MyNameSpace">
  <element name="MySecurityRootElement" type="MyPrefix:SpecificSecurity"/>
  <complexType name="SpecificSecurity">
    <sequence>
      <element name="AuthToken" type="string"/>
    </sequence>
  </complexType>
</schema>
```

A SOAP message using the XML schema and configuration above would have the following header:

```
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:SOAP-ENC="http://schemas.xmlsoap.org/soap/encoding/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <SOAP-ENV:Header>
    <m:MySecurityRootElement xmlns:m="MyNameSpace">
      <AuthToken>String</AuthToken>
    </m:MySecurityRootElement>
    ...
  </SOAP-ENV:Header>
  <SOAP-ENV:Body>
    ...
  </SOAP-ENV:Body>
```

```
</SOAP-ENV:Envelope>
```

To handle this non-default header, you must implement the method: `Directory.authenticateUserFromSOAPHeaderAPI`.

**Note**

Only available for [SOAP operations](#) [p 617].

## ***Lookup mechanism***

Because EBX® offers several authentication methods, a lookup mechanism based on conditions was set to know which method should be applied for a given request. The method application conditions are evaluated according to the authentication scheme priority. If the conditions are not satisfied, the server evaluates the next method. The following table presents the available authentication methods

for each supported protocol and their application conditions. They are ordered from the highest priority to the lowest.

| Operation / Protocol | Authentication methods and application conditions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SOAP / JMS           | <p><a href="#">SOAP Security Header</a> [p 601]</p> <ul style="list-style-type: none"> <li>• The SOAP request is received over the JMS protocol.</li> <li>• The SOAP header content must contains a Security element.</li> </ul> <p><a href="#">SOAP Specific Header</a> [p 601]</p> <ul style="list-style-type: none"> <li>• The SOAP request is received over the JMS protocol.</li> <li>• The SOAP header content must not contain a Security element.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| SOAP / HTTP          | <p><a href="#">Basic</a> [p 600]</p> <ul style="list-style-type: none"> <li>• The HTTP request must hold an Authorization header.</li> <li>• Authorization header value must start with the word Basic.</li> <li>• No login is provided in the URL parameters.</li> </ul> <p><a href="#">Standard</a> [p 600]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not hold an Authorization header.</li> <li>• A login and a password are provided in the URL parameters.</li> </ul> <p><a href="#">SOAP Security Header</a> [p 601]</p> <ul style="list-style-type: none"> <li>• The SOAP header content must contain a Security element.</li> <li>• The HTTP request must not hold an Authorization header.</li> <li>• No login is provided in the URL parameters.</li> </ul> <p><a href="#">Specific</a> [p 601]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not satisfy the conditions of the previous authentication methods.</li> </ul> <p><a href="#">SOAP Specific Header</a> [p 601]</p> <ul style="list-style-type: none"> <li>• The SOAP header content must not contain a Security element.</li> <li>• The HTTP request must not hold an Authorization header.</li> <li>• No login is provided in the URL parameters.</li> </ul> |
| WSDL / HTTP          | <p><a href="#">Basic</a> [p 600]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not hold an Authorization header.</li> <li>• Authorization header value must start with the word Basic.</li> <li>• No login is provided in the URL parameters.</li> </ul> <p><a href="#">Standard</a> [p 600]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not hold an Authorization header.</li> <li>• A login and a password are provided in the URL parameters.</li> </ul> <p><a href="#">Specific</a> [p 601]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not satisfy the conditions of the previous authentication methods.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

In case of multiple authentication methods present in the same request, EBX® will return an HTTP code 401 Unauthorized.

## 95.5 Monitoring

Data service events can be monitored through the log category `ebx.dataServices`, as declared in the EBX® main configuration file. For example, `ebx.log4j.category.log.dataServices= INFO, ebxFile:dataservices`.

### See also

[\*Configuring the EBX® logs\*](#) [p 354]

[\*TIBCO EBX® main configuration file\*](#) [p 349]

## 95.6 SOAP and REST comparative

| Operations                                                            | SOAP | REST |
|-----------------------------------------------------------------------|------|------|
| <b>Data</b>                                                           |      |      |
| Select or count records (with filter and/or view publication)         | X    | X    |
| Selector for possible enumeration values (with filter)                |      | X    |
| Insert, update or delete records                                      | X    | X    |
| Select or count history records (with filter and/or view publication) |      | X    |
| Select node values from dataset                                       | X    | X    |
| Update node value from dataset                                        |      | X    |
| Get table or dataset changes between dataspace or snapshots           | X    |      |
| Refresh a replication unit                                            | X    |      |
| Get credentials for records                                           | X    |      |
| <b>Dataspaces</b>                                                     |      |      |
| Create, close, merge a dataspace                                      | X    |      |
| Create, close a snapshot                                              | X    |      |
| Validate a dataspace or a snapshot                                    | X    |      |
| Validate a dataset                                                    | X    |      |
| Locking a dataspace                                                   | X    |      |
| <b>Workflow</b>                                                       |      |      |
| Start, resume or end a workflow                                       | X    |      |
| <b>Administration</b>                                                 |      |      |
| Manage the default directory content 'Users', 'Roles'... tables.      | X    | X    |
| Open, close the user interface                                        | X    | X    |

| Operations                                                           | SOAP | REST  |
|----------------------------------------------------------------------|------|-------|
| Select, insert, update, delete operations for administration dataset |      | X     |
| Select the system information                                        | X    | X     |
| Other                                                                |      |       |
| Develop web services from the Java API                               |      | X (*) |

(\*) See [REST Toolkit](#) [p 721] for more information.

## 95.7 Limitations

### *Date, time & dateTime format*

Data services only support the following date and time formats:

| Type        | Format                                         | Example             |
|-------------|------------------------------------------------|---------------------|
| xs:date     | yyyy-MM-dd                                     | 2007-12-31          |
| xs:time     | HH:mm:ss or HH:mm:ss.SSS                       | 11:55:00            |
| xs:dateTime | yyyy-MM-ddTHH:mm:ss or yyyy-MM-ddTHH:mm:ss.SSS | 2007-12-31T11:55:00 |

### *SOAP naming convention*

Due to the naming convention of the data service operations, each table defined within a data model must have a unique name for the WSDL generation.





## CHAPTER 96

---

# WSDL generation

This chapter contains the following topics:

1. [Supported standard](#)
2. [Operation types](#)
3. [WSDL download methods](#)
4. [HTTP examples](#)

## 96.1 Supported standard

TIBCO EBX® generates a WSDL that complies with the [W3C Web Services Description Language 1.1](#) standard.

## 96.2 Operation types

A WSDL can be generated for different types of operations:

| Operation type | WSDL description                                                                                                                                                                                                                                                                                                                                                     |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| custom         | WSDL for EBX® add-ons.                                                                                                                                                                                                                                                                                                                                               |
| dataset        | WSDL for dataset and replication operations.                                                                                                                                                                                                                                                                                                                         |
| directory      | WSDL for default EBX® directory operations. It is also possible to filter data using the <a href="#">tablePaths</a> [p 613] or <a href="#">operations</a> [p 613] parameters.                                                                                                                                                                                        |
| repository     | WSDL for dataspace or snapshot management operations.                                                                                                                                                                                                                                                                                                                |
| tables         | WSDL for operations on the tables of a specific dataset.                                                                                                                                                                                                                                                                                                             |
| userInterface  | <p>Deprecated since version 5.8.1. This operation type has been replaced by administration. While the user interface management operations are still available for backward compatibility reasons, it is recommended to no longer use this type.</p> <p>WSDL for user interface management operations (these operations can only be accessed by administrators).</p> |
| administration | <p>WSDL for administration operations like:</p> <ul style="list-style-type: none"> <li>• user interface management</li> <li>• system information retrieval</li> </ul> <p>These operations can only be accessed by administrators.</p>                                                                                                                                |
| workflow       | WSDL for EBX® workflow management operations.                                                                                                                                                                                                                                                                                                                        |

## 96.3 WSDL download methods

EBX® supports the following methods:

- from the user interface
- from HTTP protocol

A WSDL can only be downloaded by authorized profiles:

| Operation type | Access right permissions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| custom         | All profiles, if at least one web service is registered.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| dataset        | All profiles.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| directory      | All profiles, if the following conditions are valid: <ul style="list-style-type: none"> <li>No specific directory implementation is used. (The built-in Administrator role is only subject to this condition).</li> <li>Global access permissions are defined for the administration.</li> <li>'Directory' dataset permissions have writing access for the current profile.</li> </ul>                                                                                                                                                                    |
| repository     | All profiles.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| tables         | All profiles.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| userInterface  | Deprecated since version 5.8.1. This operation type has been replaced by administration. While the user interface management operations are still available for backward compatibility reasons, it is recommended to no longer use this type.<br>Built-in administrator role or delegated administrator profiles, if all conditions are valid: <ul style="list-style-type: none"> <li>Global access permissions are defined for the administration.</li> <li>'User interface' dataset permissions have writing access for the current profile.</li> </ul> |
| administration | Built-in administrator role or delegated administrator profiles, if all conditions are valid: <ul style="list-style-type: none"> <li>Global access permissions are defined for the administration.</li> <li>'Administration' dataset permissions have write access for the current profile.</li> </ul>                                                                                                                                                                                                                                                    |
| workflow       | All profiles.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## WSDL download from the user interface

An authorized user can download an EBX® WSDL from the data services administration area.

### Note

See [Generating a WSDL for dataspace operations](#) [p 182] in the user guide for more information.

## WSDL download from HTTP protocol

An application can download an EBX® WSDL using GET or POST HTTP method. The application has to be authenticated using a profile with appropriate rights.

### URL format

```
http[s]://<host>[:<port>]/<ebx-dataservices>/{type}/{dataspace}/{dataset}}?
{queryParameters}
```

Where:

- <ebx-dataservices> corresponds to the 'ebx-dataservices.war' web application's path. The path is composed by multiple, or none, URI segments followed by the web application's name.
- {type} corresponds to the [operation type](#) [p 610].
- {dataspace} corresponds to the dataspace or snapshot identifier.
- {dataset} corresponds to the dataset name.
- {queryParameters} corresponds to common or dedicated operation parameters passed through the URL.

## Parameters

A request parameter can be specified by one of the following methods:

- a PathParam which corresponds to a path segment from the URL (recommended)

- a QueryParam which corresponds to a standard HTTP parameter with value.

| Parameter name    | PathParam | QueryParam | Required | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|-----------|------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WSDL              | no        | yes        | yes      | Specifies the WSDL download operation.<br>Empty value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| login             | no        | yes        | no       | Specifies the user identifier.<br>Required when the standard authentication method is used.<br>String type value.                                                                                                                                                                                                                                                                                                                                                                                                      |
| password          | no        | yes        | no       | Specifies the user password.<br>Required when the standard authentication method is used.<br>String type value.                                                                                                                                                                                                                                                                                                                                                                                                        |
| type              | yes       | no         | yes      | Specifies the <a href="#">operation type</a> [p 610].<br>Possible values are: custom, dataset, directory, administration, userInterface, repository, tables or workflow.<br>String type value.                                                                                                                                                                                                                                                                                                                         |
| branch<br>version | yes       | yes        | (*)      | Specifies the dataspace or the snapshot identifier.<br>(*) required for tables and dataset types, ignored otherwise.<br>String type value.                                                                                                                                                                                                                                                                                                                                                                             |
| instance          | yes       | yes        | (*)      | Specifies the dataset name.<br>String type value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| tablePaths        | no        | yes        | no       | Specifies the list of table paths.<br>Optional for tables or directory types, ignored otherwise.<br>If not defined, all tables are selected.<br>Each table path is separated by a comma character.<br>String type value.                                                                                                                                                                                                                                                                                               |
| operations        | no        | yes        | no       | Allows generating a WSDL for an operations subset.<br>Optional for tables or directory operation types, ignored otherwise. If not defined, all operations for the given type are generated.<br>This parameter's value is a concatenation of one or more of the following characters: <ul style="list-style-type: none"> <li>• C = Count record(s)</li> <li>• D = Delete record(s)</li> <li>• E = Get credentials</li> <li>• G = Get changes</li> <li>• I = Insert record(s)</li> <li>• U = Update record(s)</li> </ul> |

| Parameter name        | PathParam | QueryParam | Required | Description                                                                                                                                                                                         |
|-----------------------|-----------|------------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                       |           |            |          | <ul style="list-style-type: none"> <li>• R = Read operations (equivalent to CEGS)</li> <li>• S = Select record(s)</li> <li>• W = Write operations (equivalent to DIU)</li> </ul> String type value. |
| namespaceURI          | yes       | yes        | (**)     | <p>Specifies the unique name space URI of the custom web service.</p> <p>(**)Is required when type parameter is set to custom, ignored otherwise.</p> URI type value.                               |
| attachmentFilename    | no        | yes        | (***)    | <p>Specifies the attachment file name.</p> <p>(***) optional if isContentInAttachment parameter is defined and set to true, ignored otherwise.</p> String type value.                               |
| isContentInAttachment | no        | yes        | no       | <p>Specifies if the WSDL is downloaded as an attachment.</p> Boolean type value. <p>Default value is false.</p>                                                                                     |
| targetNamespace       | no        | yes        | no       | <p>Overrides the target namespace URI of the WSDL.</p> URI type value. <p>Default value is urn:ebx:ebx-dataservices.</p>                                                                            |

### Message body

No message body is required.

## HTTP codes

An HTTP code is always returned. Errors are indicated by a code above 400.

| Status code                       | Information                                                                                                                                                                                                                                                                                                    |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 ( <i>OK</i> )                 | The WSDL content was successfully generated and is returned by the request (optionally in an <a href="#">attachment</a> [p 614]).                                                                                                                                                                              |
| 400 ( <i>Bad request</i> )        | <p>The request is incorrect. This occurs when:</p> <ul style="list-style-type: none"> <li>A request element is incorrect.</li> <li>The unicity check on table names contains at least one error.</li> </ul> <p><b>Note</b><br/>See <a href="#">WSDL and table operations</a> [p 548] for more information.</p> |
| 401 ( <i>Unauthorized</i> )       | Request requires an authenticated user.                                                                                                                                                                                                                                                                        |
| 403 ( <i>Forbidden</i> )          | Request is not allowed for the authenticated user.                                                                                                                                                                                                                                                             |
| 405 ( <i>Method not allowed</i> ) | Request is not allowed in this configuration.                                                                                                                                                                                                                                                                  |
| 500 ( <i>Internal error</i> )     | Request generates an error (a stack trace and a detailed error message are returned).                                                                                                                                                                                                                          |

## Response body

The response body depends on the returned status code and on the requested WSDL content.

- 200 (*OK*): the HTTP header Content-Type is set to text/xml; charset=UTF-8.  
If the content is in attachment, the HTTP header Content-Disposition is set to attachment; filename\*=UTF-8' '<filename.wsdl>.
- 4xx: A detailed message is returned in the body. The HTTP header Content-Type is set to text/html; charset=utf-8.

## 96.4 HTTP examples

Some of the following examples are displayed in two methods: PathParam and QueryParam.

- The WSDL will contain all repository operations, using standard authentication method.  
http[s]://<host>[:<port>]/<ebx-dataservices>/repository?WSDL&login=<login>&password=<password>
- The WSDL will contain all workflow operations.  
http[s]://<host>[:<port>]/<ebx-dataservices>/workflow?WSDL
- The WSDL will contain all tables operations for the 'dataset1' dataset in 'dataspace1' dataspace.  
PathParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/tables/dataspace1/dataset1?WSDL
```

QueryParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/tables?
```

```
WSDL&branch=dataspace1&instance=dataset1
```

- The WSDL will contain all tables with only read operations for the 'dataset1' dataset in 'dataspace1' dataspace.

PathParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/tables/dataspace1/dataset1?
```

```
WSDL&operations=R
```

QueryParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/tables?
```

```
WSDL&branch=dataspace1&instance=dataset1&operations=R
```

- The WSDL will contain the two selected tables operations for the 'dataset1' dataset in 'dataspace1' dataspace.

PathParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/tables/dataspace1/dataset1?
```

```
WSDL&tablePaths=/root/table1,/root/table2
```

QueryParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/tables?
```

```
WSDL&branch=dataspace1&instance=dataset1&tablePaths=/root/table1,/root/table2
```

- The WSDL will contain custom web service operations for the dedicated URI.

PathParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/custom/urn:ebx-
```

```
test.com.orchestranetworks.dataservices.WSDemo?WSDL
```

QueryParam

```
http[s]://<host>[:<port>]/<ebx-dataservices>/custom?WSDL&namespaceURI=urn:ebx-
```

```
test.com.orchestranetworks.dataservices.WSDemo
```



## CHAPTER 97

# SOAP operations

This chapter contains the following topics:

1. [Operations generated from a data model](#)
2. [Operations on datasets and dataspace](#)
3. [Operations on data workflows](#)
4. [Administrative services](#)

## 97.1 Operations generated from a data model

For a data model used in an TIBCO EBX® repository, it is possible to dynamically generate a corresponding WSDL, that defines its operations. When using this WSDL, it will be possible to read and/or write in the EBX® repository. For example, for a table located at the path /root/xx/exampleTable, the generated requests would follow the structure of its underlying data model and include the name of the table `<m:{operation}_exampleTable xmlns:m="urn:ebx-schemas:dataservices_1.0">`.

**Attention**

Since the WSDL and the SOAP operations tightly depend on the data model structure, it is important to redistribute the up-to-date WSDL after any data model change.

***Content policy***

Access to the content of records, the presence or absence of XML elements, depend on the [resolved permissions](#) [p 273] of the authenticated user session. Additional aspects, detailed below, can impact the content.

**Disabling fields from data model**

The `hiddenInDataServices` property, defined in the data model, allows always hiding fields in data services, regardless of the user profile. This parameter has an impact on the generated WSDL: any hidden field or group will be absent from the request and response structure.

Modifying the `hiddenInDataServices` parameter value has the following impact on a client which would still use the former WSDL:

- On request, if the data model property has been changed to `true`, and if the concerned field is present in the WSDL request, an exception will be thrown.

- On response, if the schema property has been changed to `false`, WSDL validation will return an error if it is activated.

This setting of "Default view" is defined inside data model.

#### See also

[\*Hiding a field in Data Services\*](#) [p 542]

[\*Permissions\*](#) [p 273]

## Association field

Read-access on table records can export the association fields as displayed in UI Manager. This feature can be coupled with the 'hiddenInDataServices' model parameter.

#### Note

Limitations: change and update operations do not manage association fields. Also, the select operation only exports the first level of association elements (the content of associated objects cannot contain association elements).

## Common request parameters

Several parameters are common to several operations and are detailed below.

| Element                     | Description                                                                                                                                                                                                                                                                                                                              | Required                                                                                                                       |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| branch                      | The identifier of the dataspace to which the dataset belongs.                                                                                                                                                                                                                                                                            | Either this parameter or the 'version' parameter must be defined. Required for the 'insert', 'update' and 'delete' operations. |
| version                     | The identifier of the snapshot to which the dataset belongs.                                                                                                                                                                                                                                                                             | Either this parameter or the 'branch' parameter must be defined                                                                |
| instance                    | The unique name of the dataset which contains the table to query.                                                                                                                                                                                                                                                                        | Yes                                                                                                                            |
| predicate                   | <a href="#">XPath predicate</a> [p 225] defines the records on which the request is applied. If empty, all records will be retrieved, except for the 'delete' operation where this field is mandatory.                                                                                                                                   | Only required for the 'delete' operation                                                                                       |
| data                        | Contains the records to be modified, represented within the structure of their data model. The whole operation is equivalent to an XML import. The details of the operations performed on data content are specified in the section <a href="#">Import</a> [p 213].                                                                      | Only required for the insert and update operations                                                                             |
| viewPublication             | This parameter can be combined with the <a href="#">predicate</a> [p 619] parameter as a logical AND operation.<br><br>The behavior of this parameter is described in the section <a href="#">EBX® as a Web Component</a> [p 194].<br><br>It cannot be used if the 'viewId' parameter is used, and cannot be used on hierarchical views. | No                                                                                                                             |
| viewId                      | <i>Deprecated since version 5.2.3.</i> This parameter has been replaced by the parameter 'viewPublication'. While it remains available for backward compatibility, it will eventually be removed in a future version.<br><br>This parameter cannot be used if the 'viewPublication' parameter is used.                                   | No                                                                                                                             |
| blockingConstraintsDisabled | This property is available for all table updates data service operations.<br><br>If true, the validation process disables blocking constraints defined in the data model.<br><br>If this parameter is not present, the default is false.<br><br>See <a href="#">Blocking and non-blocking constraints</a> [p 524] for more information.  | No                                                                                                                             |
| details                     | The details element specifies the following option:<br><br>The optional attribute locale (default 'en-US') defines the language of the <a href="#">blockingConstraintsDisabled</a> [p 619] parameter in which the validation messages must be returned.                                                                                  | No                                                                                                                             |

| Element                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Required |
|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| disableRedirectionToLastBroadcast | <p>This property is available for all data service operations.</p> <p>If true, access to a delivery dataspace on a D3 primary node is not redirected to the last broadcast snapshot. Otherwise, access to such a dataspace is always redirected to the last snapshot broadcast.</p> <p>If this parameter is not present, the default is false (redirection on a D3 master enabled), unless the configuration property <a href="#">ebx.dataservices.disableRedirectionToLastBroadcast.default</a> [p 357] has been set.</p> <p>If the specified dataspace is not a delivery dataspace on a D3 primary node, this parameter is ignored.</p> | No       |

## Select operations

### Select request on table

```
<m:select_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <predicate>String</predicate>
  <viewPublication>String</viewPublication>
  <exportCredentials>boolean</exportCredentials>
  <pagination>
    <previousPageLastRecordPredicate>String</previousPageLastRecordPredicate>
    <pageSize>Integer</pageSize>
  </pagination>
</m:select_{TableName}>
```

with:

| Element                                                               | Description                                                                                                                                                                                                                                                                                                                                                                   | Required                        |
|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| branch                                                                | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                          |                                 |
| version                                                               | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                          |                                 |
| instance                                                              | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                          |                                 |
| predicate                                                             | See the description under <a href="#">Common parameters</a> [p 619].<br>This parameter can be combined with the <a href="#">viewPublication</a> [p 619] parameter as a logical AND operation.                                                                                                                                                                                 |                                 |
| viewPublication                                                       | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                          |                                 |
| includesTechnicalData                                                 | The response will contain technical data if true. See also the <a href="#">optimistic locking</a> [p 636] section.<br>Each returned record will contain additional attributes for this technical information, for instance:<br><pre>...&lt;table ebxd:lastTime="2010-06-28T10:10:31.046" ebxd:lastUser="Uadmin" ebxd:uuid="9E7D0530-828C-11DF-B733-0012D01B6E76"&gt;...</pre> | No                              |
| exportCredentials                                                     | If true the select will also return the credentials for each record.                                                                                                                                                                                                                                                                                                          | No                              |
| pagination                                                            | Enables pagination, see child elements below.                                                                                                                                                                                                                                                                                                                                 | No                              |
| pageSize (nested under the pagination element)                        | When pagination is enabled, defines the number of records to retrieve.                                                                                                                                                                                                                                                                                                        | When pagination is enabled, yes |
| previousPageLastRecordPredicate (nested under the pagination element) | When pagination is enabled, XPath predicate that defines the record after which the page must fetched, this value is provided by the previous response, as the element lastRecordPredicate. If the passed record is not found, the first page will be returned.                                                                                                               | No                              |
| disableRedirectionToLastBroadcast                                     | See the description under <a href="#">Common parameters</a> [p 620].                                                                                                                                                                                                                                                                                                          |                                 |

## Select response on table

```
<ns1:select_{TableName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <data>
    <XX>
      <TableName>
        <a>key1</a>
        <b>valueb</b>
        <c>1</c>
        <d>1</d>
      </TableName>
    </XX>
  </data>
  <credentials>
    <XX>
      <TableName predicate="./a='key1'">
        <a>W</a>
        <b>W</b>
      </TableName>
    </XX>
  </credentials>
</ns1:select_{TableName}Response>
```

```
<c>W</c>
<d>W</d>
</TableName>
</XX>
</credentials>
<lastRecordPredicate>./a='key1'</lastRecordPredicate>
</ns1:select_{TableName}Response>
```

with:

| Element             | Description                                                                                                                                                                      |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| data                | Content of records that are displayed following the table path.                                                                                                                  |
| credentials         | Contains the access right for each node of each record.                                                                                                                          |
| lastRecordPredicate | Only returned if the pagination is enabled, this defines the last records in order to be used on the next request in the element <code>previousPageLastRecordPredicated</code> . |

See also the [optimistic locking](#) [p 636] section.

### Select request on dataset

This operation returns dataset content without table.

```
<m:selectInstance xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
</m:selectInstance>
```

with:

| Element                           | Description                                                          | Required |
|-----------------------------------|----------------------------------------------------------------------|----------|
| branch                            | See the description under <a href="#">Common parameters</a> [p 619]. |          |
| version                           | See the description under <a href="#">Common parameters</a> [p 619]. |          |
| instance                          | See the description under <a href="#">Common parameters</a> [p 619]. |          |
| disableRedirectionToLastBroadcast | See the description under <a href="#">Common parameters</a> [p 620]. |          |

### Select response on dataset

```
<ns1:selectInstanceResponse xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <data>
    <settings>
      <XX>
        <a>key1</a>
        <b>valueb</b>
        <c>1</c>
        <d>true</d>
      </XX>
    </settings>
  </data>
</ns1:selectInstanceResponse>
```

with:

| Element | Description                    |
|---------|--------------------------------|
| data    | Dataset content without table. |

## Delete operation

Deletes records or, for a child dataset, defines the record state as "occulting" or "inherited" according to the record context. Records are selected by the predicate parameter.

### Delete request

```
<m:delete_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <predicate>String</predicate>
  <includeOcculting>boolean</includeOcculting>
  <inheritIfInOccultingMode>boolean</inheritIfInOccultingMode>
  <checkNotChangedSinceLastTime>dateTime</checkNotChangedSinceLastTime>
  <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>
  <details locale="Locale"/>
</m:delete_{TableName}>
```

with:

| Element                           | Description                                                                                                                                          | Required |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| branch                            | See the description under <a href="#">Common parameters</a> [p 619].                                                                                 |          |
| instance                          | See the description under <a href="#">Common parameters</a> [p 619].                                                                                 |          |
| predicate                         | See the description under <a href="#">Common parameters</a> [p 619].                                                                                 |          |
| includeOcculting                  | Includes the records in occulting mode.<br>Default value is false.                                                                                   | No       |
| inheritIfInOccultingMode          | Inherits the record from its parent if it is in occulting mode.<br>Default value is false.                                                           | No       |
| occultIfInherit                   | <i>Deprecated since version 5.7.0</i> Occults the record if it is in inherit mode.<br>Default value is false.                                        | No       |
| checkNotChangedSinceLastTime      | Timestamp used to ensure that the record has not been modified since the last read. Also see the <a href="#">optimistic locking</a> [p 636] section. | No       |
| blockingConstraintsDisabled       | See the description under <a href="#">Common parameters</a> [p 619].                                                                                 |          |
| details                           | See the description under <a href="#">Common parameters</a> [p 619].                                                                                 |          |
| disableRedirectionToLastBroadcast | See the description under <a href="#">Common parameters</a> [p 620].                                                                                 |          |

## Delete response

If one of the provided parameters is illegal, if a required parameter is missing, if the action is not authorized or if no record is selected, an exception is returned. Otherwise, the specific response is returned:

```
<ns1:delete_{TableName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
  <blockingConstraintMessage>String</blockingConstraintMessage>
</ns1:delete_{TableName}Response>
```

with:

| Element                   | Description                                                                                                                                                                                           |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| status                    | '00' indicates that the operation has been executed successfully.<br>'95' indicates that at least one operation has violated a blocking constraint, resulting in the overall operation being aborted. |
| blockingConstraintMessage | This element is present if the status is equal to '95' with a localized message. The locale of the message is retrieved from the request parameter or from the user session.                          |

## Count operation

### Count request

```
<m:count_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <predicate>String</predicate>
</m:count_{TableName}>
```

with:

| Element                           | Description                                                          |
|-----------------------------------|----------------------------------------------------------------------|
| branch                            | See the description under <a href="#">Common parameters</a> [p 619]. |
| version                           | See the description under <a href="#">Common parameters</a> [p 619]. |
| instance                          | See the description under <a href="#">Common parameters</a> [p 619]. |
| predicate                         | See the description under <a href="#">Common parameters</a> [p 619]. |
| disableRedirectionToLastBroadcast | See the description under <a href="#">Common parameters</a> [p 620]. |

### Count response

```
<ns1:count_{TableName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <count>Integer</count>
</ns1:count_{TableName}Response>
```



with:

| Element | Description                                                            |
|---------|------------------------------------------------------------------------|
| count   | The number of records that correspond to the predicate in the request. |

## Update operation

### Update request

```
<m:update_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <updateOrInsert>boolean</updateOrInsert>
  <byDelta>boolean</byDelta>
  <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>
  <details locale="Locale"/>
  <data>
    <XX>
      <TableName>
        <a>String</a>
        <b>String</b>
        <c>String</c>
        <d>String</d>
        ...
      </TableName>
    </XX>
  </data>
</m:update_{TableName}>
```

with:

| Element                           | Description                                                                                                                                                                                                                                                                                                   | Required |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| branch                            | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                          |          |
| instance                          | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                          |          |
| updateOrInsert                    | If true and the record does not currently exist, the operation creates the record.<br>boolean type, the default value is false.                                                                                                                                                                               | No       |
| byDelta                           | If true and an element does not currently exist in the incoming message, the target value is not changed.<br>If false and node is declared hiddenInDataServices, the target value is not changed.<br>The complete behavior is described in the sections <a href="#">Insert and update operations</a> [p 214]. | No       |
| blockingConstraintsDisabled       | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                          |          |
| details                           | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                          |          |
| disableRedirectionToLastBroadcast | See the description under <a href="#">Common parameters</a> [p 620].                                                                                                                                                                                                                                          |          |
| data                              | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                          |          |

See also [Optimistic locking](#) [p 636]

## Update response

```
<ns1:update_{TableName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
  <blockingConstraintMessage>String</blockingConstraintMessage>
</ns1:update_{TableName}Response>
```

with:

| Element                   | Description                                                                                                                                                                                           |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| status                    | '00' indicates that the operation has been executed successfully.<br>'95' indicates that at least one operation has violated a blocking constraint, resulting in the overall operation being aborted. |
| blockingConstraintMessage | This element is present if the status is equal to '95' with a localized message. The locale of the message is retrieved from the request parameter or from the user session.                          |

## Insert operation

### Insert request

```
<m:insert_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <byDelta>boolean</byDelta>
  <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>
  <details locale="Locale"/>
  <data>
    <XX>
      <TableName>
        <a>String</a>
        <b>String</b>
        <c>String</c>
        <d>String</d>
        ...
      </TableName>
    </XX>
  </data>
</m:insert_{TableName}>
```

with:

| Element                           | Description                                                                                                                                                                                                                                                                                                                  | Required |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| branch                            | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                         |          |
| instance                          | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                         |          |
| byDelta                           | <p>If true and an element does not currently exist in the incoming message, the target value is not changed.</p> <p>If false and node is declared hiddenInDataServices, the target value is not changed.</p> <p>The complete behavior is described in the sections <a href="#">Insert and update operations</a> [p 214].</p> | No       |
| blockingConstraintsDisabled       | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                         |          |
| details                           | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                         |          |
| disableRedirectionToLastBroadcast | See the description under <a href="#">Common parameters</a> [p 620].                                                                                                                                                                                                                                                         |          |
| data                              | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                         |          |

### Insert response

```
<ns1:insert_{TableName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
  <blockingConstraintMessage>String</blockingConstraintMessage>
  <inserted>
    <predicate>./a='String'</predicate>
  </inserted>
</ns1:insert_{TableName}Response>
```

with:

| Element                   | Description                                                                                                                                                                                           |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| status                    | '00' indicates that the operation has been executed successfully.<br>'95' indicates that at least one operation has violated a blocking constraint, resulting in the overall operation being aborted. |
| blockingConstraintMessage | This element is present if the status is equal to '95' with a localized message. The locale of the message is retrieved from the request parameter or from the user session.                          |
| predicate                 | A predicate matching the primary key of the inserted record. When several records are inserted, the predicates follow the declaration order of the records in the input message.                      |

## Get changes operations

Returns changes according to the [Content policy](#) [p 617].

## Get changes requests

### Changes between two datasets:

```
<m:getChangesOnDataSet_{schemaName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <compareWithBranch>String</compareWithBranch>
  <compareWithVersion>String</compareWithVersion>
  <compareWithInstance>String</compareWithInstance>
  <resolvedMode>boolean</resolvedMode>
  <includeInstanceUpdates>boolean</includeInstanceUpdates>
</m:getChangesOnDataSet_{schemaName}>
```

### Changes between two tables:

```
<m:getChanges_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <compareWithBranch>String</compareWithBranch>
  <compareWithVersion>String</compareWithVersion>
  <resolvedMode>boolean</resolvedMode>
</m:getChanges_{TableName}>
```

with:

| Element                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Required                                                                                                     |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| branch                                     | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                              |
| version                                    | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                              |
| instance                                   | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                              |
| compareWithBranch                          | The identifier of the dataspace with which to compare.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | One of either this parameter or the ' <a href="#">compareWithVersion</a> [p 629]' parameter must be defined. |
| compareWithVersion                         | The identifier of the snapshot with which to compare.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | One of either this parameter or the ' <a href="#">compareWithBranch</a> [p 629]' parameter must be defined.  |
| compareWithInstance                        | The identifier of the dataset with which to compare. If it is undefined, <a href="#">instance</a> [p 629] parameter is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | No                                                                                                           |
| resolvedMode                               | Defines whether or not the difference is calculated in resolved mode. Default is true.<br>See <b>Resolved mode</b> <code>DifferenceHelper.resolvedMode</code> <sup>opt</sup> for more information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | No                                                                                                           |
| includeInstanceUpdates                     | Defines if the content updates of the dataset are included. Default is false.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | No                                                                                                           |
| pagination                                 | Enables pagination context for the operations <code>getChanges</code> and <code>getChangesOnDataSet</code> .<br><br>Allows client to define pagination context size. Each page contains a collection of inserted, updated and/or deleted records of tables according to the maximum size.<br><br>Get changes persisted context is built at first call according to the page size parameter in request.<br><br>The pagination context is persisted <a href="#">on the server file system</a> [p 382] and allows invoking the next page until last page or when a timeout is reached.<br><br>For creation: Defines <code>pageSize</code> parameter.<br><br>For next: Defines context element with <code>identifier</code> from previous response.<br><br>Enables pagination, see child elements below. | No                                                                                                           |
| pageSize (nested under pagination element) | Defines maximum number of records in each page. Minimal size is 50.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | No (Only for creation)                                                                                       |
| context (nested under pagination element)  | Defines content of pagination context.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No (Only for next)                                                                                           |

| Element                                   | Description                                                          | Required |
|-------------------------------------------|----------------------------------------------------------------------|----------|
| identifier (nested under context element) | Pagination context identifier.                                       | Yes      |
| disableRedirectionToLastBroadcast         | See the description under <a href="#">Common parameters</a> [p 620]. |          |

**Note**

If none of the *compareWithBranch* or *compareWithVersion* parameters are specified, the comparison will be made with their parent:

- if the current dataspace or snapshot is a dataspace, the comparison is made with its initial snapshot (includes all changes made in the dataspace);
- if the current dataspace or snapshot is a snapshot, the comparison is made with its parent dataspace (includes all changes made in the parent dataspace since the current snapshot was created);
- returns an exception if the current dataspace is the 'Reference' dataspace.

See also *DifferenceHelper*<sup>API</sup>

## Get changes responses

### Changes between two datasets:

```
<ns1:getChangesOnDataSet_{schemaName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <updated>
    <changes>
      <path>... Path of changed terminal value ...</path>
      <path>...</path>
    </changes>
    <data>
      ... see the whole content of dataset values (without table) ...
    </data>
  </updated>
  <getChanges_{TableName1}>
    ... see the getChanges between tables response example ...
  </getChanges_{TableName1}>
  <getChanges_{TableName2}>
    ... see the getChanges between tables response example ...
  </getChanges_{TableName2}>
  ...
</ns1:getChangesOnDataSet_{schemaName}Response>
```

### Changes between two tables:

```
<ns1:getChanges_{TableName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <inserted>
    <XX>
      <TableName>
        <a>AVALUE3</a>
        <b>BVALUE3</b>
        <c>CVALUE3</c>
        <d>DVALUE3</d>
      </TableName>
    </XX>
  </inserted>
  <updated>
    <changes>
      <change predicate="./a='AVALUE2'">
        <path>/b</path>
        <path>/c</path>
      </change>
    </changes>
    <data>
      <XX>
        <TableName>
          <a>AVALUE2</a>
```

```
<b>BVALUE2.1</b>
<c>CVALUE2.1</c>
<d>DVALUE2</d>
</TableName>
</XX>
</data>
</updated>
<deleted>
  <predicate>./a='AVALUE1'</predicate>
</deleted>
</ns1:getChanges_{TableName}Response>
```

with:

| Element                                   | Description                                                                                                                                                                                                                                                                                                                                                                   | Required                     |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| inserted                                  | Contains inserted record(s) from choice <code>compareWithBranch</code> or <code>compareWithVersion</code> .<br>Content under this element corresponding to an XML export of inserted records.                                                                                                                                                                                 | No                           |
| updated                                   | Contains updated record(s) or dataset content.                                                                                                                                                                                                                                                                                                                                | No                           |
| changes (nested under updated element)    | Only the group of field have been updated.                                                                                                                                                                                                                                                                                                                                    | Yes                          |
| change (nested under changes element)     | Group of fields have been updated with own XPath predicate attribute of the record.                                                                                                                                                                                                                                                                                           | Yes                          |
| path (nested under change element)        | Path in the record.                                                                                                                                                                                                                                                                                                                                                           | Yes                          |
| path (nested under changes element)       | Path in the dataset.                                                                                                                                                                                                                                                                                                                                                          | Yes                          |
| data (nested under updated element)       | Content under this element corresponding to an XML export of dataset or updated records.                                                                                                                                                                                                                                                                                      | No                           |
| deleted                                   | Records have been deleted from context of request.<br>Content corresponding to a list of predicate element who contains the XPath predicate of record.                                                                                                                                                                                                                        | No                           |
| pagination                                | When pagination is enabled on request.<br>Get changes persisted context allows invoking the next page until last page or when the context timeout is reached.<br>Contains a next page: Defines context element with <code>identifier</code> .<br>Is the last page: Defines context element without <code>identifier</code> .<br>Enables pagination, see child elements below. | No                           |
| context (nested under pagination element) | Defines content of pagination context.                                                                                                                                                                                                                                                                                                                                        | Yes (Only for next and last) |
| identifier (nested under context element) | Pagination context identifier. Not defined at last returned page.                                                                                                                                                                                                                                                                                                             | No                           |
| pageNumber (nested under context element) | Current page number in pagination context.                                                                                                                                                                                                                                                                                                                                    | Yes                          |
| totalPages (nested under context element) | Total pages in pagination context.                                                                                                                                                                                                                                                                                                                                            | Yes                          |

## Get changes operation with pagination enabled

Only pagination element and sub elements have been described.



For creation:

Extract of request:

```
...
<pagination>
  <!-- on first request for creation -->
  <pageSize>Integer</pageSize>
</pagination>
...
```

Extract of response:

```
...
<pagination>
  <!-- on next request to continue -->
  <context>
    <identifier>String</identifier>
    <pageNumber>Integer</pageNumber>
    <totalPages>Integer</totalPages>
  </context>
</pagination>
...
```

For next:

Extract of request:

```
...
<pagination>
  <context>
    <identifier>String</identifier>
  </context>
</pagination>
...
```

Extract of response:

```
...
<pagination>
  <!-- on next request to continue -->
  <context>
    <identifier>String</identifier>
    <pageNumber>Integer</pageNumber>
    <totalPages>Integer</totalPages>
  </context>
</pagination>
...
```

For last:

Extract of request:

```
...
<pagination>
  <context>
    <identifier>String</identifier>
  </context>
</pagination>
...
```

Extract of response:

```
...
<pagination>
  <context>
    <pageNumber>Integer</pageNumber>
    <totalPages>Integer</totalPages>
  </context>
</pagination>
...
```

# Get credentials operation

## Get credentials request

```
<m:getCredentials_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <predicate>String</predicate>
  <viewPublication>String</viewPublication>
</m:getCredentials_{TableName}>
```

with:

| Element         | Description                                                          | Required |
|-----------------|----------------------------------------------------------------------|----------|
| branch          | See the description under <a href="#">Common parameters</a> [p 619]. |          |
| version         | See the description under <a href="#">Common parameters</a> [p 619]. |          |
| instance        | See the description under <a href="#">Common parameters</a> [p 619]. |          |
| predicate       | See the description under <a href="#">Common parameters</a> [p 619]. |          |
| viewPublication | See the description under <a href="#">Common parameters</a> [p 619]. |          |

## Get credentials response

```
<ns1:getCredentials_{TableName}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <XX>
    <TableName>
      <a>R</a>
      <b>W</b>
      <c>H</c>
      <d>W</d>
      ...
    </TableName>
  </XX>
</ns1:getCredentials_{TableName}Response>
```

With the following possible values:

- R: for read-only
- W: for read-write
- H: for hidden

# Multiple chained operations

## Multiple operations request

It is possible to run multiple operations across tables in the dataset, while ensuring a consistent response. The operations are executed sequentially, according to the order defined on the client side. All operations are executed in a single transaction with a `SERIALIZABLE` isolation level. If all requests in the multiple operation are read-only, they are allowed to run fully concurrently along with other read-only transactions, even in the same dataspace.

When an error occurs during one operation in the sequence, all updates are rolled back and the client receives a `StandardException` error message with details.

See [Concurrency and isolation levels](#) [p 468].

```
<m:multi_ xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>
  <details locale="Locale"/>
  <request id="id1">
    <{operation}__{TableName}>
      ...
    </{operation}__{TableName}>
  </request>
  <request id="id2">
    <{operation}__{TableName}>
      ...
    </{operation}__{TableName}>
  </request>
</m:multi_>
```

with:

| Element                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                 | Required |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| branch                            | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                        |          |
| version                           | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                        |          |
| instance                          | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                        |          |
| blockingConstraintsDisabled       | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                        |          |
| details                           | See the description under <a href="#">Common parameters</a> [p 619].                                                                                                                                                                                                                                                                                                                                                        |          |
| disableRedirectionToLastBroadcast | See the description under <a href="#">Common parameters</a> [p 620].                                                                                                                                                                                                                                                                                                                                                        |          |
| request                           | <p>This element contains one operation, like a single operation without branch, version and instance parameters. This element can be repeated multiple times for additional operations. Each request can be identified by an 'id' attribute. In a response, this 'id' attribute is returned for identification purposes.</p> <p>Operations such as count, select, getChanges, getCredentials, insert, delete or update.</p> | Yes      |

**Note:**

- Does not accept a limit on the number of request elements.
- The request id attribute must be unique in multi-operation requests.
- If all operations are read only (count, select, getChanges, or getCredentials) then the whole transaction is set as read-only for performance considerations.

**Limitation:**

- The multi operation applies to one model and one dataset (parameter instance).
- The select operation cannot use the pagination parameter.

See also

*Procedure*<sup>API</sup>

*Repository*<sup>API</sup>

Multiple operations response

See each response operation for details.

```
<ns1:multi_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <response id="id1">
    <{operation}_{TableName}Response>
      ...
    </{operation}_{TableName}Response>
  </response>
  <response id="id2">
    <{operation}_{TableName}Response>
      ...
    </{operation}_{TableName}Response>
  </response>
</ns1:multi_Response>
```

with:

| Element  | Description                                                                                                                                                                                                                                                                                                                                                                               |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| response | <p>This element contains the response of one operation. It is be repeated multiple times for additional operations. Each response is identified by an 'id' attribute set in the request or automatically generated.</p> <p>The content of the element corresponds to the response of a single operation, such as count, select, getChanges, getCredentials, insert, delete or update.</p> |

Optimistic locking

To prevent an update or a delete operation on a record that was read earlier but may have changed in the meantime, an optimistic locking mechanism is provided.

A select request can include technical information by adding the element `includesTechnicalData`:

```
<m:select_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <predicate>String</predicate>
  <includesTechnicalData>boolean</includesTechnicalData>
</m:select_{TableName}>
```

The value of the `lastTime` attribute can then be used in the following update request. If the record has been changed since the specified time, the update will be cancelled. The attribute `lastTime` has to be added on the record to prevent the update of a modified record.

```
<m:update_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <updateOrInsert>true</updateOrInsert>
  <data>
    <XX>
      <TableName ebxd:lastTime="2010-06-28T10:10:31.046">
        <a>String</a>
        <b>String</b>
        <c>String</c>
        <d>String</d>
        ...
      </TableName>
    </XX>
  </data>
```

```
</m:update_{TableName}>
```

The value of the `lastTime` attribute can also be used to prevent deletion on a modified record:

```
<m:delete_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">  
  <branch>String</branch>  
  <version>String</version>  
  <instance>String</instance>  
  <predicate>String</predicate>  
  <checkNotChangedSinceLastTime>2010-06-28T10:10:31.046</checkNotChangedSinceLastTime>  
</m:delete_{TableName}>
```

#### Note

The element `checkNotChangedSinceLastTime` may be used more than once but only for the same record. This implies that if the predicate element returns more than one record, the request will fail.

## 97.2 Operations on datasets and dataspace

Parameters for operations on dataspace and snapshots are as follows:

| <i>Element</i>             | <i>Description</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <i>Required</i>                                                                                                                                                   |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| branch                     | Identifier of the target dataspace on which the operation is applied. When not specified, the 'Reference' dataspace is used except for the merge dataspace operation where it is required.                                                                                                                                                                                                                                                                                                                                                                                                      | One of either this parameter or the 'version' parameter must be defined. Required for the dataspace merge, locking, unlocking and replication refresh operations. |
| version                    | Identifier of the target snapshot on which the operation is applied.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | One of either this parameter or the 'branch' parameter must be defined                                                                                            |
| versionName                | Identifier of the snapshot to create. If empty, it will be defined on the server side.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | No                                                                                                                                                                |
| childBranchName            | Identifier of the dataspace child to create. If empty, it will be defined on the server side.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | No                                                                                                                                                                |
| instance                   | The unique name of the dataset on which the operation is applied.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Required for the replication refresh operation.                                                                                                                   |
| ensureActivation           | Defines if validation must also check whether this instance is activated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Yes                                                                                                                                                               |
| details                    | <p>Defines if validation returns details.</p> <p>The optional attribute <code>severityThreshold</code> defines the lowest severity level of message to return. The possible values are, in descending order of severity, 'fatal', 'error', 'warning', or 'info'. For example, setting the value to 'error' will return error and fatal validation messages. If this attribute is not defined, all levels of messages are returned by default.</p> <p>The optional attribute <code>locale</code> (default 'en-US') defines the language in which the validation messages are to be returned.</p> | No. If not specified, no details are returned.                                                                                                                    |
| owner                      | <p>Defines the owner.</p> <p>Must respect the inner format as returned by <code>Profile.format<sup>API</sup></code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | No                                                                                                                                                                |
| branchToCopyPermissionFrom | Defines the identifier of the dataspace from which to copy the permissions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | No                                                                                                                                                                |
| documentation              | Documentation for a dedicated language.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | No                                                                                                                                                                |

| <i>Element</i>                                       | <i>Description</i>                                                 | <i>Required</i>                                      |
|------------------------------------------------------|--------------------------------------------------------------------|------------------------------------------------------|
|                                                      | Multiple documentation elements may be used for several languages. |                                                      |
| locale (nested under the documentation element)      | Locale of the documentation.                                       | Only required when the documentation element is used |
| label (nested under the documentation element)       | Label for the language.                                            | No                                                   |
| description (nested under the documentation element) | Description for the language.                                      | No                                                   |

## Validate a dataspace

### Validate dataspace request

```
<m:validate xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
</m:validate>
```

### Validate dataspace response

```
<ns1:validate_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <validationReport>
    <instanceName>String</instanceName>
    <fatals>boolean</fatals>
    <errors>boolean</errors>
    <infos>boolean</infos>
    <warnings>boolean</warnings>
  </validationReport>
</ns1:validate_Response>
```

## Validate a dataset

### Validate dataset request

```
<m:validateInstance xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <version>String</version>
  <instance>String</instance>
  <ensureActivation>boolean</ensureActivation>
  <details severityThreshold="fatal|error|warning|info" locale="Locale"/>
</m:validateInstance>
```

### Validate dataset response

```
<ns1:validateInstance_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <validationReport>
    <instanceName>String</instanceName>
    <fatals>boolean</fatals>
    <errors>boolean</errors>
    <infos>boolean</infos>
    <warnings>boolean</warnings>
    <details>
      <reportItem>
        <severity>{fatal|error|warning|info}</severity>
        <message>
          <internalId />
          <text>String</text>
        </message>
        <subject>
          <table>Path</table>
          <predicate>String</predicate>
        </subject>
      </reportItem>
    </details>
  </validationReport>
</ns1:validateInstance_Response>
```

```
<path>Path</path>
</subject>
</reportItem>
</details>
</validationReport>
</ns1:validateInstance_Response>
```

## Create a dataspace

### Create dataspace request

```
<m:createBranch xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <owner>String</owner>
  <branchToCopyPermissionFrom>String</branchToCopyPermissionFrom>
  <documentation>
    <locale>Locale</locale>
    <label>String</label>
    <description>String</description>
  </documentation>
  <childBranchName>String</childBranchName>
</m:createBranch>
```

### Create dataspace response

```
<ns1:createBranch_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
  <childBranchName>String</childBranchName>
</ns1:createBranch_Response>
```

with:

| Element | Description                                                       |
|---------|-------------------------------------------------------------------|
| status  | '00' indicates that the operation has been executed successfully. |

## Create a snapshot

### Create snapshot request

```
<m:createVersion xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <versionName>String</versionName>
  <owner>String</owner>
  <documentation>
    <locale>Locale</locale>
    <label>String</label>
    <description>String</description>
  </documentation>
</m:createVersion>
```

### Create snapshot response

```
<ns1:createVersion_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
  <versionName>String</versionName>
</ns1:createVersion_Response>
```

with:

| Element | Description                                                       |
|---------|-------------------------------------------------------------------|
| status  | '00' indicates that the operation has been executed successfully. |



## Locking a dataspace

### Lock dataspace request

```
<m:lockBranch xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <durationToWaitForLock>Integer</durationToWaitForLock>
  <message>
    <locale>Locale</locale>
    <label>String</label>
  </message>
</m:lockBranch>
```

with:

| Element                                   | Description                                                                                                   | Required                                       |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------|------------------------------------------------|
| durationToWaitForLock                     | This parameter defines the maximum duration (in seconds) that the operation waits for a lock before aborting. | No, does not wait by default                   |
| message                                   | User message of the lock. Multiple message elements may be used.                                              | No                                             |
| locale (nested under the message element) | Locale of the user message.                                                                                   | Only required when the message element is used |
| label (nested under the message element)  | The user message.                                                                                             | No                                             |

### Lock dataspace response

```
<ns1:lockBranch_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:lockBranch_Response>
```

with:

| Element | Description                                                                                                                                                                               |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| status  | '00' indicates that the operation has been executed successfully.<br>'94' indicates that the dataspace has been already locked by another user.<br>Otherwise, a SOAP exception is thrown. |

## Unlocking a dataspace

### Unlock dataspace request

```
<m:unlockBranch xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
</m:unlockBranch>
```

## Unlock datasource response

```
<ns1:unlockBranch_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:unlockBranch_Response>
```

with:

| Element | Description                                                                                                 |
|---------|-------------------------------------------------------------------------------------------------------------|
| status  | '00' indicates that the operation has been executed successfully.<br>Otherwise, a SOAP exception is thrown. |

## Merge a datasource

### Merge datasource request

```
<m:mergeBranch xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <deleteDataOnMerge>boolean</deleteDataOnMerge>
  <deleteHistoryOnMerge>boolean</deleteHistoryOnMerge>
</m:mergeBranch>
```

with:

| Element              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Required |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| deleteDataOnMerge    | This parameter is available for the merge datasource operation. Sets whether the specified datasource and its associated snapshots will be deleted upon merge.<br><br>When this parameter is not specified in the request, the default value is false. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.dataDeletionOnCloseOrMerge.default</code> in the <a href="#">EBX® main configuration file</a> [p 357].<br><br>See <a href="#">Deleting data and history</a> [p 380] for more information.                           | No       |
| deleteHistoryOnMerge | This parameter is available for the merge datasource operation. Sets whether the history associated with the specified datasource will be deleted upon merge. Default value is false.<br><br>When this parameter is not specified in the request, the default value is false. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.historyDeletionOnCloseOrMerge.default</code> in the <a href="#">EBX® main configuration file</a> [p 357].<br><br>See <a href="#">Deleting data and history</a> [p 380] for more information. | No       |

#### Note

The merge decision step is bypassed during merges performed through data services. In such cases, the data in the child datasource automatically overrides the data in the parent datasource.

### Merge datasource response

```
<ns1:mergeBranch_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:mergeBranch_Response>
```

with:

| Element | Description                                                       |
|---------|-------------------------------------------------------------------|
| status  | '00' indicates that the operation has been executed successfully. |

## Close a dataspace or snapshot

### Close dataspace or snapshot request

#### Close dataspace request:

```
<m:closeBranch xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <deleteDataOnClose>boolean</deleteDataOnClose>
  <deleteHistoryOnClose>boolean</deleteHistoryOnClose>
</m:closeBranch>
```

#### Close snapshot request:

```
<m:closeVersion xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <version>String</version>
  <deleteDataOnClose>boolean</deleteDataOnClose>
</m:closeVersion>
```

with:

| Element              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Required |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| deleteDataOnClose    | <p>This parameter is available for the close dataspace and close snapshot operations. Sets whether the specified snapshot, or dataspace and its associated snapshots, will be deleted upon closure.</p> <p>When this parameter is not specified in the request, the default value is false. It is possible to redefine this default value by specifying the property <code>ebx.dataservices.dataDeletionOnCloseOrMerge.default</code> in the <a href="#">EBX® main configuration file</a> [p 357].</p> <p>See <a href="#">Deleting data and history</a> [p 380] for more information.</p> | No       |
| deleteHistoryOnClose | <p>This parameter is available for the close dataspace operation. Sets whether the history associated with the specified dataspace will be deleted upon closure. Default value is false.</p> <p>When this parameter is not specified in the request, the default value is false. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.historyDeletionOnCloseOrMerge.default</code> in the <a href="#">EBX® main configuration file</a> [p 357].</p> <p>See <a href="#">Deleting data and history</a> [p 380] for more information.</p>          | No       |

### Close dataspace or snapshot response

#### Close dataspace response:

```
<ns1:closeBranch_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:closeBranch_Response>
```

#### Close snapshot request:

```
<ns1:closeVersion_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:closeVersion_Response>
```

# Replication refresh

## Replication refresh request

```
<m:replicationRefresh_${schema} xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>String</branch>
  <instance>String</instance>
  <unitName>String</unitName>
</m:replicationRefresh_${schema}>
```

with:

| Element  | Description                                                                                           | Required |
|----------|-------------------------------------------------------------------------------------------------------|----------|
| branch   | See the description under <a href="#">Common parameters</a> [p 638].                                  | Yes      |
| instance | See the description under <a href="#">Common parameters</a> [p 638].                                  | Yes      |
| unitName | Name of the replication unit.<br><br>See also <a href="#">Replication refresh information</a> [p 258] | Yes      |

## Replication refresh response

```
<ns1:replicationRefresh_${schema}Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:replicationRefresh_${schema}Response>
```

with:

| Element | Description                                                       |
|---------|-------------------------------------------------------------------|
| status  | '00' indicates that the operation has been executed successfully. |

## 97.3 Operations on data workflows

Parameters for data workflows operations are retrieved from the SOAP header in the session.  
*Deprecated since version 5.7.0 to define parameters in the SOAP message body.*

See [session parameters](#) [p 599] for more information.

| Element                                                                                   | Description                                                                                                                                                                                                                                                                  | Required |
|-------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| parameters                                                                                | <p><i>Deprecated since version 5.7.0</i> While it remains available for backward compatibility, it will eventually be removed in a future major version.</p> <p><b>Note</b></p> <p>The parameters element is ignored if at least one session parameter has been defined.</p> | No       |
| parameter (nested under the parameters element). Multiple parameter elements may be used. | An input parameter for the workflow.                                                                                                                                                                                                                                         | No       |
| name (nested under the parameter element)                                                 | Name of the parameter.                                                                                                                                                                                                                                                       | Yes      |
| value (nested under the parameter element)                                                | Value of the parameter.                                                                                                                                                                                                                                                      | No       |

## Start a workflow

Start a workflow from a workflow launcher. It is possible to start a workflow with localized documentation and specific input parameters (with name and optional value).

### Note

The workflow creator is initialized from the session and the workflow priority is retrieved from the last published version.

### Sample request:

```
<m:workflowProcessInstanceStart xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <publishedProcessKey>String</publishedProcessKey>
  <documentation>
    <locale>Locale</locale>
    <label>String</label>
    <description>String</description>
  </documentation>
</m:workflowProcessInstanceStart>
```

with:

| Element             | Description                                                                                                | Required |
|---------------------|------------------------------------------------------------------------------------------------------------|----------|
| publishedProcessKey | Identifier of the workflow launcher.                                                                       | Yes      |
| documentation       | See the description under <a href="#">Common parameters</a> [p 638].                                       | No       |
| parameters          | <i>Deprecated since version 5.7.0</i> See the description under <a href="#">Common parameters</a> [p 645]. | No       |

### Sample response:

```
<m:workflowProcessInstanceStart_Response xmlns:m="urn:ebx-schemas:dataservices_1.0">
```

```
<processInstanceKey>String</processInstanceKey>
</m:workflowProcessInstanceStart_Response>
```

with:

| Element            | Description                                                                                                                                                                                                                   | Required |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| processInstanceId  | <i>Deprecated since version 5.6.1</i> This parameter has been replaced by the 'processInstanceKey' parameter. While it remains available for backward compatibility, it will eventually be removed in a future major version. | No       |
| processInstanceKey | Workflow identifier.                                                                                                                                                                                                          | No       |

## Resume a workflow

Resume a workflow in a wait step from a resume identifier. It is possible to define specific input parameters (with name and optional value).

### Sample request:

```
<m:workflowProcessInstanceResume xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <resumeId>String</resumeId>
</m:workflowProcessInstanceResume>
```

with:

| Element    | Description                                                                                                | Required |
|------------|------------------------------------------------------------------------------------------------------------|----------|
| resumeId   | Resume identifier of the waiting task.                                                                     | Yes      |
| parameters | <i>Deprecated since version 5.7.0</i> See the description under <a href="#">Common parameters</a> [p 645]. | No       |

### Sample response:

```
<m:workflowProcessInstanceResume_Response xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
  <processInstanceKey>String</processInstanceKey>
</m:workflowProcessInstanceResume_Response>
```

with:

| Element            | Description                                                                                                                                                                           | Required |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| status             | '00' indicates that the operation has been executed successfully.<br>'20' indicates that the workflow has not been found.<br>'21' indicates that the event has already been received. | Yes      |
| processInstanceKey | Identifier of the workflow. This parameter is returned if the operation has been executed successfully.                                                                               | No       |

## End a workflow

End a workflow from its identifier.

### Sample request:

```
<m:workflowProcessInstanceEnd xmlns:m="urn:ebx-schemas:dataservices_1.0">
```

```
<processInstanceKey>String</processInstanceKey>
</m:workflowProcessInstanceEnd>
```

with:

| Element             | Description                                                                                                                                                                                                                                       | Required                                                                                           |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| processInstanceKey  | Identifier of the workflow.                                                                                                                                                                                                                       | Either this parameter or 'publishedProcessKey' and 'processInstanceId' parameters must be defined. |
| publishedProcessKey | <i>Deprecated since version 5.6.1</i> Due to a limitation this parameter has been replaced by the 'processInstanceKey' parameter. While it remains available for backward compatibility, it will eventually be removed in a future major version. | No                                                                                                 |
| processInstanceId   | <i>Deprecated since version 5.6.1</i> Due to a limitation this parameter has been replaced by the 'processInstanceKey' parameter. While it remains available for backward compatibility, it will eventually be removed in a future major version. | No                                                                                                 |

#### Sample response:

```
<m:workflowProcessInstanceEnd_Response xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</m:workflowProcessInstanceEnd_Response>
```

with:

| Element | Description                                                       | Required |
|---------|-------------------------------------------------------------------|----------|
| status  | '00' indicates that the operation has been executed successfully. | Yes      |

## 97.4 Administrative services

### Directory services

The services on directory provide operations on the 'Users' and 'Roles' tables of the default directory. To execute an operation related to these services, the authenticated user must be a member of the built-in role 'Administrator'.

The technical dataspace and dataset must be set to ebx-directory. For all SOAP operation syntaxes, see [Operations generated from a data model](#) [p 617] for more information.

### Create a user in the directory

This example of a SOAP insert request adds a user to the EBX® directory.

```
<m:insert_user xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <branch>ebx-directory</branch>
  <instance>ebx-directory</instance>
  <data>
    <directory>
      <users>
        <login>login</login>
        <lastName>lastname</lastName>
```

```

<firstName>firstname</firstName>
<email>firstname.lastname@email.com</email>
<password>***</password>
<passwordMustChange>true</passwordMustChange>
<builtInRoles>
  <administrator>false</administrator>
  <readOnly>false</readOnly>
</builtInRoles>
<comments>a comment</comments>
</users>
</directory>
</data>
</m:insert_user>

```

For the insert SOAP response syntax, see [insert response](#) [p 627] for more information.

## User interface operations

See [Application locking](#) [p 390] for more information.

Parameters for operations on the user interface are as follows:

| Element       | Description                                                                   | Required |
|---------------|-------------------------------------------------------------------------------|----------|
| closedMessage | Message to be displayed to users when the user interface is closed to access. | No       |

### Close user interface request

The close operation removes all user sessions that are not acceptable in this mode.

```

<m:close xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <closedMessage>Access is temporarily forbidden.</closedMessage>
</m:close>

```

### Close user interface response

```

<ns1:close_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:close_Response>

```

### Open user interface request

```

<m:open xmlns:m="urn:ebx-schemas:dataservices_1.0"/>

```

### Open user interface response

```

<ns1:open_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <status>String</status>
</ns1:open_Response>

```

## System information operation

This operation returns the EBX® system information. The information returned is the same as the information contained in the log header `kernel.log` or in the UI tab 'Administration' > 'System Information'. The response contains several keys, labels, and values representing the configuration and status of EBX®. To execute this operation, the authenticated user must be a member of the built-in role 'Administrator'.



## Parameters

The following parameter is applicable.

| Parameter | Description                                                                                                                                                                                                           | Required                                                                  |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| details   | <p>Defines attributes that must be applied to response messages.</p> <p>The attribute <code>locale</code> (default: EBX® default locale) defines the language in which the system item messages must be returned.</p> | No, but if specified, the <code>locale</code> attribute must be provided. |

## System information request

This SOAP request will return all EBX® instance's system information and format them using "en\_US" locale.

```
<m:systemInformation xmlns:m="urn:ebx-schemas:dataservices_1.0">
  <details locale="en_US" />
</m:systemInformation>
```

## System information response

```
<ns1:systemInformation_Response xmlns:ns1="urn:ebx-schemas:dataservices_1.0">
  <bootInfoEBX>
    <label>String</label>
    <infoItem>
      <key>String</key>
      <label>String</label>
      <content>String</content>
      <content>String</content>
      ...
    </infoItem>
    ...
  </bootInfoEBX>
  <repositoryInfo>
    <label>String</label>
    <infoItem>
      <key>String</key>
      <label>String</label>
      <content>String</content>
      <content>String</content>
      ...
    </infoItem>
    ...
  </repositoryInfo>
  <bootInfoVM>
    <label>String</label>
    <infoItem>
      <key>String</key>
      <label>String</label>
      <content>String</content>
      ...
    </infoItem>
    ...
  </bootInfoVM>
</ns1:systemInformation_Response>
```



---

# REST data services

---

## CHAPTER 98

# Introduction

This chapter contains the following topics:

1. [Overview](#)
2. [Activation and configuration](#)
3. [Interactions](#)
4. [Security](#)
5. [Monitoring](#)
6. [SOAP and REST comparative](#)
7. [Limitations](#)

## 98.1 Overview

REST data services allow external systems to interact with data governed in the TIBCO EBX® repository using the RESTful built-in services.

The request and response syntax for built-in services are described in the chapter [Built-in RESTful services](#) [p 659].

Built-in REST data services allow to perform operations such as:

- Selecting, inserting, updating, deleting, or counting records
- Selecting or counting history records
- Selecting, updating, or counting dataset values
- Administrative operations to manage access to the UI or to system information

**Note**

See [SOAP and REST comparative](#) [p 657].

## 98.2 Activation and configuration

REST and SOAP Data services are activated by deploying the `ebx-dataservices` web application along with the other EBX® modules. See [Java EE deployment overview](#) [p 315] for more information.

In case of specific deployment, for example using reverse-proxy mode, see [URLs computing](#) [p 359] for more information.

Currently only protocol HTTP(S) is supported.

## 98.3 Interactions

### *Input and output message encoding*

All input and output messages must be *exclusively* in UTF-8 for REST built-in.

### *Tracking information*

Depending on the data services operation being called, it may be possible to specify session tracking information.

- Example for a RESTful operation, the JSON request contains:

```
{
  "procedureContext": // JSON Object (optional)
  {
    "trackingInformation": "String" // JSON String (optional)
  }, ...
}
```

For more information, see `Session.getTrackingInfoAPI` in the Java API.

### *Session parameters*

Depending on the data services operation being called, it is possible to specify session input parameters. They are defined in the request body.

Input parameters are available on custom Java components with a session object, such as: triggers, access rules, custom web services. They are also available on data workflow operations.

- Example for a RESTful operation, the JSON request contains:

```
{
  "procedureContext": // JSON Object (optional)
  {
    "trackingInformation": "String", // JSON String (optional)
    "inputParameters": // JSON Array (optional)
    [
      // JSON Object for each parameter
      {
        "name": "String" // JSON String (required)
        "value": "String" // JSON String (optional)
      },
      ...
    ]
  }, ...
}
```

For more information, see `Session.getInputParameterValueAPI` in the Java API.

### *Exception handling*

When an error occurs, a JSON exception response is returned to the caller. For example:

```
{
  "code": 999, // JSON Number, HTTP status code
  "errors": [
    {
      "severity": "...", // JSON String, severity (optional)
      "rowIndex": 999, // JSON Number, request row index (optional)
      "userCode": "...", // JSON String, user code (optional)
      "message": "...", // JSON String, message
      "details": "...", // JSON String, URL (optional)
      "pathInRecord": "...", // JSON String, Path in record (optional)
      "pathInDataset": "...", // JSON String, Path in dataset (optional)
    }
  ]
}
```

```
}

```

The response contains an HTTP status code and a table of errors. The severity of each error is specified by a character, with one of the possible values (F: fatal, E: error, W: warning, I: information).

The HTTP error 422 (*Unprocessable entity*) corresponds to a functional error. It contains a user code under the `userCode` key and is a JSON String type.

See also [HTTP codes](#) [p 665]

See also `Severity.toParsableString`<sup>API</sup>

## 98.4 Security

### Authentication

Authentication is mandatory to access built-in services. Several authentication methods are available and described below. The descriptions are ordered by priority (EBX® applies the highest priority authentication method first).

- 'Token Authentication Scheme' method is based on the HTTP-Header Authorization, as described in [RFC 2617](#).

```
> Authorization: <tokenType> <accessToken>

```

For more information on this authentication scheme, see [Token authentication operations](#) [p 668].

See also [HTTP Authorization header policy](#) [p 359]

- 'Basic Authentication Scheme' method is based on the HTTP-Header Authorization in base 64 encoding, as described in [RFC 2617 \(Basic Authentication Scheme\)](#).

```
If the user agent wishes to send the userid "Alibaba" and password "open sesame",
it will use the following header field:
> Authorization: Basic QWxpymFiYTpvGvUHNlc2FtZQ==

```

#### Note

The [WWW-Authenticate](#) [p 664] header can be valued with this method.

See also [HTTP Authorization header policy](#) [p 359]

- 'Standard Authentication Scheme' is based on the HTTP Request. User and password are extracted from request parameters. For more information on request parameters, see [Parameters](#) [p 612] section.

For more information on this authentication scheme, see `Directory.authenticateUserFromLoginPassword`<sup>API</sup>.

- The 'REST Forward Authentication Scheme' is used only when calling a REST service from a [user service](#) [p 562], that reuses the current authenticated session.

For more information, see [Implementing a user service](#) [p 573] making a call to [REST data services](#) [p 580].

- 'Specific authentication Scheme' is based on the HTTP Request. For example, an implementation can extract a password-digest or a ticket from the HTTP Request. See `Directory.authenticateUserFromHttpRequest`<sup>API</sup> for more information.

- 'Anonymous authentication Scheme' is used only to access the REST services handling the authentication operations. The credentials acquisition, password changes, etc. imply that the user cannot be known yet.

## Global permissions

Global access permissions can be independently defined for the REST built-in services access. For more information see [Global permissions](#) [p 385].

## Lookup mechanism

Because EBX® offers several authentication methods, a lookup mechanism based on conditions was set to know which method should be applied for a given request. The method application conditions are evaluated according to the authentication scheme priority. If the conditions are not satisfied, the server evaluates the next method. The following table presents the available authentication methods for each supported protocol and their application conditions. They are ordered from the highest priority to the lowest.

| Operation / Protocol | Authentication methods and application conditions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| REST / HTTP          | <p><a href="#">Token</a> [p 654]</p> <ul style="list-style-type: none"> <li>• The HTTP request must hold an Authorization header.</li> <li>• Authorization header value must start with the word EBX.</li> <li>• No login is provided in the URL parameters.</li> </ul> <p><a href="#">Basic</a> [p 654]</p> <ul style="list-style-type: none"> <li>• The HTTP request must hold an Authorization header.</li> <li>• Authorization header value must start with the word Basic.</li> <li>• No login is provided in the URL parameters.</li> </ul> <p><a href="#">Standard</a> [p 654]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not hold an Authorization header.</li> <li>• A login and a password are provided in the URL parameters.</li> </ul> <p><a href="#">Rest forward</a> [p 654]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not contain an Authorization header.</li> <li>• No login is provided in the URL parameters.</li> </ul> <p><a href="#">Specific</a> [p 654]</p> <ul style="list-style-type: none"> <li>• The HTTP request must not satisfy the conditions of the previous authentication methods.</li> </ul> <p><a href="#">Anonymous</a> [p 655]</p> <ul style="list-style-type: none"> <li>• None of the previous authentication methods can be applied.</li> <li>• The requested REST service is handling an authentication operation.</li> </ul> |

In case of multiple authentication methods present in the same request, EBX® will return an HTTP code 401 Unauthorized.

## 98.5 Monitoring

Data service events can be monitored through the log category `ebx.dataServices`, as declared in the EBX® main configuration file. For example, `ebx.log4j.category.log.dataServices= INFO, ebxFile:dataservices`.

### See also

[\*Configuring the EBX® logs\*](#) [p 354]

[\*TIBCO EBX® main configuration file\*](#) [p 349]



## 98.6 SOAP and REST comparative

| Operations                                                            | SOAP | REST |
|-----------------------------------------------------------------------|------|------|
| <b>Data</b>                                                           |      |      |
| Select or count records (with filter and/or view publication)         | X    | X    |
| Selector for possible enumeration values (with filter)                |      | X    |
| Insert, update or delete records                                      | X    | X    |
| Select or count history records (with filter and/or view publication) |      | X    |
| Select node values from dataset                                       | X    | X    |
| Update node value from dataset                                        |      | X    |
| Get table or dataset changes between dataspace or snapshots           | X    |      |
| Refresh a replication unit                                            | X    |      |
| Get credentials for records                                           | X    |      |
| <b>Dataspaces</b>                                                     |      |      |
| Create, close, merge a dataspace                                      | X    |      |
| Create, close a snapshot                                              | X    |      |
| Validate a dataspace or a snapshot                                    | X    |      |
| Validate a dataset                                                    | X    |      |
| Locking a dataspace                                                   | X    |      |
| <b>Workflow</b>                                                       |      |      |
| Start, resume or end a workflow                                       | X    |      |
| <b>Administration</b>                                                 |      |      |
| Manage the default directory content 'Users', 'Roles'... tables.      | X    | X    |
| Open, close the user interface                                        | X    | X    |

| Operations                                                           | SOAP | REST  |
|----------------------------------------------------------------------|------|-------|
| Select, insert, update, delete operations for administration dataset |      | X     |
| Select the system information                                        | X    | X     |
| Other                                                                |      |       |
| Develop web services from the Java API                               |      | X (*) |

(\*) See [REST Toolkit](#) [p 721] for more information.

## 98.7 Limitations

### *Date, time & dateTime format*

Data services only support the following date and time formats:

| Type        | Format                                         | Example             |
|-------------|------------------------------------------------|---------------------|
| xs:date     | yyyy-MM-dd                                     | 2007-12-31          |
| xs:time     | HH:mm:ss or HH:mm:ss.SSS                       | 11:55:00            |
| xs:dateTime | yyyy-MM-ddTHH:mm:ss or yyyy-MM-ddTHH:mm:ss.SSS | 2007-12-31T11:55:00 |

### *JMS*

- JMS protocol is not supported.

## CHAPTER 99

# Built-in RESTful services

This chapter contains the following topics:

1. [Introduction](#)
2. [Request](#)
3. [Response](#)
4. [Administration operations](#)
5. [Token authentication operations](#)
6. [Data operations](#)
7. [Limitations](#)

## 99.1 Introduction

The architecture used is called ROA (Resource-Oriented Architecture), it can be an alternative to SOA (Service-Oriented Architecture). The chosen resources are readable and/or writable by third-party systems, according to the request content.

The HATEOAS approach of the built-in RESTful services also allows to experience an intuitive and straightforward navigation, which implies that the data details could be obtained through a link.

**Note**

All operations are stateless.

## 99.2 Request

This chapter describes the elements to use in order to build a conform REST request, such as: the HTTP method, the URL format, the header fields and the message body.

**See also**

[Interactions](#) [p 653]

[Security](#) [p 654]

### ***HTTP method***

Considered HTTP methods for built-in RESTful services, are:

- GET: used to select master data defined in the URL (the URL size limit depends on the application server or on the browser, that must be lower than or equal to 2KB).
- POST: used to insert one or more records in a table or to select the master data defined in the URL (the size limit is 2MB or more depending on the application server. Each parameter is limited to a value containing 1024 characters).
- PUT: used to update the master data defined in the URL.
- DELETE: used to delete either the record defined in the URL or multiple records defined with the table URL and the record table in the message body.

## URL

REST URL contains:

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/{categoryVersion}/
{specificPath}[:{extendedAction}]?{queryParameters}
```

Where:

- <ebx-dataservices> corresponds to the 'ebx-dataservices.war' web application's path. The path is composed by multiple, or none, URI segments followed by the web application's name.
- {category} corresponds to the [operation category](#) [p 661].
- {categoryVersion} corresponds to the category version: current value is v1.
- {specificPath} corresponds to a specific path inside the category.
- {extendedAction} corresponds to the extended action name (optional).
- {queryParameters} corresponds to [common](#) [p 663] or dedicated operation parameters passed by the URL.

## Operation category

It specializes the operation, it is added in the path of the URL in {category} and it takes one of the following values:

|                |                                                                                                                                                                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>admin</b>   | Administration operations reserved to administrators.<br>For more information, see: <a href="#">Administration operations</a> [p 666].                                                                                             |
| <b>auth</b>    | Manage token authentication method.<br>For more information, see: <a href="#">Token authentication operations</a> [p 668] and <a href="#">Token Authentication Scheme</a> [p 654].                                                 |
| <b>data</b>    | Lists dataset content, requests a table, a record or a field record content, including modified operations on dataset node, table, record and record field.<br>For more information, see: <a href="#">Data operations</a> [p 671]. |
| <b>history</b> | Lists history dataset content, requests a history table, a history of a record or a history record.<br>For more information, see: <a href="#">Data operations</a> [p 671].<br><br>See also <a href="#">History</a> [p 249]         |

## Header fields

These header field definitions are used by TIBCO EBX®.

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Accept</b>           | <p>Used to specify content types by order of preference to be used in the response, the first supported one will be chosen and specified in the response header Content-Type. Currently, the only supported one is application/json. If none is supported, the result depends on the <code>ebx.dataservices.rest.request.checkAccept</code> property:</p> <ul style="list-style-type: none"> <li>• If true, an HTTP error response is returned with code 406.</li> <li>• If false, the response is returned with the default content type, that is application/json.</li> </ul> <p>See also <a href="#">Configuring data services</a> [p 357]</p> |
| <b>Accept-Language</b>  | <p>Used for specifying the preferred locale for the response. The supported locales are defined in the schema model.</p> <p>If none of the preferred locale are supported, the default locale for the current model is used.</p>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Authorization</b>    | <p>Used for 'Basic Authentication Scheme' and 'Token Authentication Scheme' methods, otherwise the request is rejected.</p> <p>See also <a href="#">Authentication</a> [p 654]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Content-Type</b>     | <p>Used to specify the request body media type. The supported types are application/json and application/x-www-form-urlencoded. The request value is checked and if it is not supported, then an HTTP error message is returned with the code 415 (<i>Unsupported media type</i>).</p> <p>See also <a href="#">Configuring data services</a> [p 357]</p>                                                                                                                                                                                                                                                                                          |
| <b>X-Requested-With</b> | <p>If present and in case of authentication failure, prevents the addition of the <code>www-Authenticate</code> header in the response.</p> <p>See also <a href="#">Response header <code>www-Authenticate</code></a> [p 664]</p>                                                                                                                                                                                                                                                                                                                                                                                                                 |

See [RFC2616](#) for more information about HTTP Header Field Definitions.

## Common parameters

These optional parameters are available for all data service operations.

| Parameter                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>disableRedirectionToLastBroadcast</code> | <p>This parameter only has impact on a D3 architecture.</p> <p>If <code>true</code>, access to a delivery dataspace on a D3 primary node is not redirected to the last broadcast snapshot. Otherwise, access to such a dataspace is always redirected to the last broadcast snapshot.</p> <p>If the specified dataspace is not a delivery dataspace on a D3 primary node, this parameter is ignored.</p> <p>Boolean type value. If this parameter is not present, the default is <code>false</code> (redirection to a D3 master enabled), unless the configuration property <a href="#">ebx.dataservices.disableRedirectionToLastBroadcast.default</a> [p 357] has been set.</p> <p>See also <a href="#">Primary node</a> [p 429]</p> |
| <code>indent</code>                            | <p>Specifies if the response should be indented, to be easier to read for a human.</p> <p>Boolean type, default value is <code>false</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Message body

It contains the request data using the JSON format, see [JSON Request body](#) [p 696].

### Note

Requests may define a message body only when using POST or PUT HTTP methods.

## 99.3 Response

This chapter describes the responses returned by built-in RESTful services.

- See [Exception handling](#) [p 653] for details on standard error handling (where the HTTP code is greater than or equal to 300).

## Header fields

These header field definitions are used by EBX®.

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Content-Language</b> | Indicates the locale used in the response for labels and descriptions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Content-Type</b>     | Indicates the response body content type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Location</b>         | If a new record has been successfully inserted, the query URL for this record is returned by this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>WWW-Authenticate</b> | <p>This header field is added to the HTTP response when authentication fails with the 401 (<i>Unauthorized</i>) HTTP code. Its value consists of a list with at least one authentication method applicable to the request URI. It is present if and only if the following conditions are verified:</p> <ul style="list-style-type: none"> <li>the 'Basic Authentication Scheme' method is enabled and</li> <li>the X-Requested-With HTTP header is not present.</li> </ul> <p>If the client is able to interpret the authentication method, it is possible to resubmit the request providing the appropriate credentials.</p> <p>The <code>administration</code> property <code>ebx.dataservices.rest.auth.tryBasicAuthentication</code> [p 357] must be set to <code>true</code>.</p> <p>See also</p> <p><a href="#">Request header X-Requested-With</a> [p 662]</p> <p><a href="#">Authentication</a> [p 654]</p> |



## HTTP codes

| HTTP code                             | Description                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 ( <i>OK</i> )                     | The request has been successfully handled.                                                                                                                                                                                                                                                                                                                                                                   |
| 201 ( <i>Created</i> )                | A new record has been created, in this case, the header field <code>Location</code> is returned with its resource URL.                                                                                                                                                                                                                                                                                       |
| 204 ( <i>No content</i> )             | The request has been successfully handled but no response body is returned.                                                                                                                                                                                                                                                                                                                                  |
| 400 ( <i>Bad request</i> )            | the request URL or body is not well-formed or contains invalid content.                                                                                                                                                                                                                                                                                                                                      |
| 401 ( <i>Unauthorized</i> )           | Authentication has failed.                                                                                                                                                                                                                                                                                                                                                                                   |
| 403 ( <i>Forbidden</i> )              | <p>Permission was denied to read or modify the specified resource for the authenticated user.</p> <p>This error is also returned when the user:</p> <ul style="list-style-type: none"> <li>is not allowed to modify a field mentioned in the request message body.</li> <li>is not allowed to access the REST connector.</li> </ul> <p>For more details, see <a href="#">Global permissions</a> [p 655].</p> |
| 404 ( <i>Not found</i> )              | The resource specified in the URL cannot be found.                                                                                                                                                                                                                                                                                                                                                           |
| 406 ( <i>Not acceptable</i> )         | Content type defined in the request's <code>Accept</code> parameter is not supported. This error can be returned only if the EBX® property <code>ebx.rest.request.checkAccept</code> is set to <code>true</code> .                                                                                                                                                                                           |
| 409 ( <i>Conflict</i> )               | <p>A concurrent modification has occurred.</p> <p>See also <a href="#">Optimistic locking</a> [p 691]</p>                                                                                                                                                                                                                                                                                                    |
| 415 ( <i>Unsupported media type</i> ) | The request content is not supported, the request header value <code>Content-Type</code> is not supported by the operation.                                                                                                                                                                                                                                                                                  |
| 422 ( <i>Unprocessable entity</i> )   | The new resource's content cannot be accepted for semantic reasons.                                                                                                                                                                                                                                                                                                                                          |
| 500 ( <i>Internal error</i> )         | Unexpected error thrown by the application. Error details can usually be found in EBX® logs.                                                                                                                                                                                                                                                                                                                 |

## Message body

The response body content's format depends on the HTTP code value:

- HTTP codes from 200 included to 300 excluded: the content format depends on the associated request ([JSON](#) [p 699] samples).

With the exception of code 204 (*No content*).

- HTTP codes greater than or equal to 300: the content describes the error. See [JSON](#) [p 653] for details on the format.

## 99.4 Administration operations

Administration operations are related to:

- the administration category.
- the administration dataspace accessible through the data category.

### Note

administration category and administration dataspace can only be used by administrators.

### Directory operations

The EBX® default directory configuration is manageable with built-in RESTful services. The users and roles tables, the mailing lists and other objects are concerned. For more information, see [Users and roles directory](#) [p 401].

### Note

Triggers are present on the directory's tables, ensuring the data consistency.

The URL format is:

`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/Bebx-directory/ebx-directory`

### Directory configuration operations

The EBX® default directory configuration is manageable like dataset nodes. It can be accessed and modified through the data category operations. Each field is self-described when metadata is requested.

See [select](#) [p 671] and [update](#) [p 684] operations for more information.

### Mailing lists operations

There are two default mailing lists that can be configured in the EBX® directory: one for everybody and one for administrators. These lists can be handled like dataset nodes through the data category operations.

See [select](#) [p 671] and [update](#) [p 684] operations for more information.

### Directory users operations

Users can be manipulated like records of the data category using the operations of the latter. For security purposes, an administrator cannot delete himself. The user's salutation must be chosen among the ones available in the 'salutations' table.

See [select](#) [p 671], [update](#) [p 684], [insert](#) [p 679] and [delete](#) [p 686] operations for more information.

### Directory roles operations

Roles are records of the data category and can be managed with its operations. EBX® roles are assigned to users through the 'usersRoles' association table. 'usersRoles' is automatically fed when the directory is administered through the user interface. However, it is not the case through data services and role assignments require manual operations. Roles inclusions are specified in the 'rolesInclusions'

association table. As for the 'usersRoles' table, the management of roles inclusions requires manual operations. Each table is self-descriptive when metadata is requested.

See [select](#) [p 671], [update](#) [p 684], [insert](#) [p 679] and [delete](#) [p 686] operations for more information.

## ***User interface operations***

The EBX® user interface can be opened or closed to users for maintenance needs. Handled information is similar to what is contained in the UI tab 'Administration' > 'User interface configuration' > 'Advanced perspective' > 'Graphical interface configuration' > 'Application locking'.

URL format is:

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/Bebx-manager/ebx-manager/
domain/toolStatus
```

See also [Application locking](#) [p 390]

### **Retrieve user interface state**

User interface status and the unavailability message are accessible like dataset nodes.

See [Select operation](#) [p 671] and the [JSON](#) [p 706] example, for more information.

### **Open or close user interface**

User interface status and the unavailability message can be modified like dataset nodes using the update operation. To open the user interface set the content of toolStatus to true, or to false to close it.

See [Update operation](#) [p 684] and the [JSON](#) [p 699] examples, for more information.

## ***System information operation***

This operation returns system information on the EBX® server. This is accepted for GET and POST HTTP methods. Warning: no update will be possible in the POST HTTP method because the request body is ignored. The information returned is the same as the information contained in the log header kernel.log or in the UI tab 'Administration' > 'System Information'. The response contains several keys, labels, and values representing the configuration and status of EBX®. The mode of representation of the response may be flat or hierarchical.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/admin/v1/systemInformation
```

See also

[TIBCO EBX® main configuration file](#) [p 349]

[Repository administration](#) [p 374]

## Parameters

The following parameter is applicable.

| Parameter             | Description                                                                                                                                                                                                                                                                                            |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| systemInformationMode | <p>Specifies the returned mode:</p> <ul style="list-style-type: none"> <li>flat: A flat representation under the following information groups: bootInfoEBX, repositoryInfo and bootInfoVM.</li> <li>hierarchical: A hierarchical representation.</li> </ul> <p>String type, default value is flat.</p> |

## HTTP codes

| HTTP code                  | Description                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 ( <i>OK</i> )          | The system information was successfully returned.                                                                                                                                                                                                                                                                   |
| 400 ( <i>Bad request</i> ) | <p>The request is not correct, it contains one of the following errors:</p> <ul style="list-style-type: none"> <li>the HTTP method is not GET nor POST,</li> <li>the HTTP parameter systemInformationMode is not correct,</li> <li>the operation is not supported.</li> <li>the request path is invalid.</li> </ul> |
| 403 ( <i>Forbidden</i> )   | The user is not an administrator.                                                                                                                                                                                                                                                                                   |

## Response body

It is returned, if and only if, the HTTP code is 200 (*OK*). The content structure depends on the provided parameter systemInformationMode or its default value.

See the [JSON](#) [p 700] example of the flat representation.

See the [JSON](#) [p 700] example of the hierarchical representation.

## 99.5 Token authentication operations

These operations allow to create or revoke an authentication token. Authentication tokens have a timeout period. If a token is not used to access the EBX® server within this period, it will automatically be revoked. This timeout period is refreshed on each access to EBX® server.

### Note

The token timeout is modifiable through the administration property [ebx.dataservices.rest.auth.token.timeout](#) [p 357] (the default value is 30 minutes).

### Create token operation

This operation requires using the POST HTTP method with a request containing the user's credentials and, optionally, [session parameters](#) [p 653].

URL format is:

`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/auth/v1/token:create`

## Message body

A message body must be defined in the HTTP request. It necessarily contains one of the following set of data:

- A login and a password value. Both JSON attributes are mandatory and of String types.  
See `Directory.authenticateUserFromLoginPasswordAPI` for more information.
- The specific JSON attribute set to true. When activated, this flag allows to performed a user authentication against the whole HTTP request. Warning, even if login and password attributes are defined in the JSON request's body, setting specific to true lead to a specific user authentication.

See `Directory.authenticateUserFromHttpRequestAPI` for more information.

See the [JSON](#) [p 696] examples of a token creation request.

## HTTP codes

| HTTP code                  | Description                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 (OK)                   | The token was successfully created.                                                                                                                                                                                                                                                                                                                          |
| 400 (Bad request)          | For one of the following reasons: <ul style="list-style-type: none"> <li>• the syntax is not correct,</li> <li>• the HTTP method is not POST,</li> <li>• the operation is not supported.</li> </ul>                                                                                                                                                          |
| 401 (Unauthorized)         | For one of the following reasons: <ul style="list-style-type: none"> <li>• The login and/or password is/are incorrect.</li> <li>• Authentication data for other methods is defined.</li> </ul>                                                                                                                                                               |
| 422 (Unprocessable entity) | For one of the following reasons: <ul style="list-style-type: none"> <li>• PasswordMustChange: The password must be changed (only available with the default directory). See <a href="#">Change password operation</a> [p 670].</li> <li>• RestrictedAccess: User access is closed for maintenance or other actions (reserved to administrators).</li> </ul> |

## Response body

If the HTTP code is 200 (OK), the body holds the token value and its type.

See the [JSON](#) [p 701] example of a token creation response.

The token can later be used to authenticate a user by setting the HTTP-Header Authorization accordingly.

See also ['Token authentication Scheme' method](#) [p 654]

## Change password operation

This operation modifies the password of an existing user account. It can be used in an authenticated context: `login` parameter, if present, is checked against the current session or taken from it, if absent. It could also be used in an unauthenticated context, for example when the [Create token operation](#) [p 668] aborts with the HTTP code 422 (*Unprocessable entity*) with reason: `PasswordMustChange`.

It requires the use of:

- the EBX® default directory
- the POST HTTP method
- the message body containing the structure specified below

URL format is:

`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/auth/v1/user:changePassword`

## Message body

The message body must be defined in the request. It necessarily contains a `password` and a `passwordNew`, the `login` is optional (all are `String`).

See the [JSON](#) [p 696] example of a password change and token creation request.

## HTTP codes

| HTTP code                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 204 ( <i>No content</i> )           | The password has been changed.                                                                                                                                                                                                                                                                                                                                                                                                               |
| 400 ( <i>Bad request</i> )          | For one of the following reasons: <ul style="list-style-type: none"> <li>• the EBX® default directory is required,</li> <li>• the syntax is not correct,</li> <li>• the HTTP method is not POST,</li> <li>• the provided <code>login</code> and the user's session one mismatch</li> <li>• the operation is not supported.</li> </ul>                                                                                                        |
| 401 ( <i>Unauthorized</i> )         | For the following reason: <ul style="list-style-type: none"> <li>• The <code>login</code> and/or <code>password</code> is/are incorrect.</li> </ul>                                                                                                                                                                                                                                                                                          |
| 422 ( <i>Unprocessable entity</i> ) | For one of the following reasons: <ul style="list-style-type: none"> <li>• <code>PasswordChangeAbort: passwordNew</code> is empty.</li> <li>• <code>PasswordChangeAbort: passwordNew</code> is equal to <code>password</code>.</li> <li>• <code>PasswordChangeAbort: password</code> is incorrect.</li> <li>• <code>RestrictedAccess</code>: User access is closed for maintenance or other actions (reserved to administrators).</li> </ul> |

## Response body

If HTTP code 204 (*No content*) is returned, then the password has been modified.

## Revoke token operation

This operation requires using the POST HTTP method. No message body is needed.

URL format is:

`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/auth/v1/token:revoke`

### Header fields

#### Authorization

This field is required, `tokenType` and `accessToken` fields must have the values returned from the "token create" operation.

```
> Authorization: <tokenType> <accessToken>
```

### HTTP codes

| HTTP code                   | Description                                                                                                                                                                                                                              |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 204 ( <i>No content</i> )   | The token has been revoked successfully.                                                                                                                                                                                                 |
| 400 ( <i>Bad request</i> )  | For one of the following reasons: <ul style="list-style-type: none"> <li>the configuration is not activated,</li> <li>the syntax is incorrect,</li> <li>the HTTP method is not POST,</li> <li>the operation is not supported.</li> </ul> |
| 401 ( <i>Unauthorized</i> ) | Authentication has failed.                                                                                                                                                                                                               |

## 99.6 Data operations

Operations from the data operation category concern the datasets, the dataset fields, tables, records or record fields.

Operations from the history category and concern historized content from datasets, tables, records or the record fields.

### Select operation

Select operation may use one of the following methods:

- GET HTTP method,
- POST HTTP method without message body or
- POST HTTP method with message body and optionally [session parameters](#) [p 653].

URL formats are:

- Dataset tree**, depending on [operation category](#) [p 661]:

The data category returns the hierarchy of the selected dataset, this includes group and table nodes.

The history category returns the hierarchy of the selected history dataset, this includes the pruned groups for history table nodes only.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/{dataspace}/{dataset}
```

#### Note

Terminal nodes and sub-nodes are not included.

- **Dataset node:** the data category returns the terminal nodes contained in the selected node.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/{dataset}/{pathInDataset}
```

#### Note

Not applicable with the history category.

- **Table**, depending on [operation category](#) [p 661]:

the data category returns the table content and/or metadata, current page records and URLs for pagination.

The history category returns the history table content and/or metadata, current page records and URLs for pagination.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/{dataspace}/{dataset}/{pathInDataset}
```

See also [Count operation](#) [p 688]

- **Record**, depending on [operation category](#) [p 661]:

the data category returns the record content and/or metadata.

The history category returns history record content and/or metadata.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}
```

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/{dataspace}/{dataset}/{pathInDataset}?primaryKey={xpathExpression}
```

#### Note

The record access by the primary key (primaryKey parameter) is limited to its root node. It is recommended to use the encoded primary key, available in the details field in order to override this limitation. Similarly, for a history record, use the encoded primary key, available in the historyDetails field.

- **Field**, depending on [operation category](#) [p 661]:

the data category returns the field record content where structure depends on its type.



The history category returns the field history record content where structure depends on its type.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/{dataspace}/
{dataset}/{pathInDataset}/{encodedPrimaryKey}/{pathInRecord}
```

#### Note

The field must be either an association node, a selection node, a terminal node or above.

Where:

- {category} corresponds to the [operation category](#) [p 661].
- {dataspace} corresponds to B followed by the dataspace identifier or to v followed by the snapshot identifier.
- {dataset} corresponds to the dataset identifier.
- {pathInDataset} corresponds to the path of the dataset node, that can be a group node or a table node.
- {encodedPrimaryKey} corresponds to the encoded representation of the primary key.

**See also** *RESEncodingHelper<sup>API</sup>*

- {xpathExpression} corresponds to the record primary key, using the XPath expression.
- {pathInRecord} corresponds to the path starting from the table node.

## Parameters

The following parameters are applicable to the select operation.

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| includeContent | Includes the content field with the content corresponding to the selection.<br>Boolean type, default value is true.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| includeDetails | Includes the details field in the metadata and the response, for each indirect reachable resource. The returned value corresponds to its URL resource.<br>Type Boolean, default value is true.<br><br>See also <a href="#">includeMeta</a> [p 674]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| includeHistory | Includes those fields for a historized content: <ul style="list-style-type: none"> <li>The history property in the metadata. The returned value corresponds to a boolean value.</li> <li>The historyDetails property in the content and for each indirectly reachable resource. This point is coupled to the <a href="#">includeDetails</a> [p 674] parameter. The returned value corresponds to its URL resource.</li> </ul> Boolean type, default value is false.<br><br><div> <b>Note</b><br/> The includeHistory parameter is ignored in the history category, the default value is true. </div><br>See also<br><a href="#">includeMeta</a> [p 674]<br><a href="#">includeDetails</a> [p 674]<br><i>AdaptationTable.getHistory<sup>API</sup></i> |
| includeLabel   | Includes the label field associated with each simple type content.<br>Possible values are: <ul style="list-style-type: none"> <li>yes: the label is included for the foreign key, enumeration, record and selector values.</li> <li>all: the label field is included, as for the yes value and also for the <a href="#">Content of simple type</a> [p 713].</li> </ul><br>See also <i>SchemaNode.displayOccurrence<sup>API</sup></i><br><br><ul style="list-style-type: none"> <li>no: the label field is not included (integration use case).</li> </ul> String type, default value is yes.<br><br><div> <b>Note</b><br/> The label field is not included if it is equal to the content field. </div>                                               |
| includeMeta    | Includes the meta field corresponding to the description of the structure returned in the content field.<br>Boolean type, default value is false.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     | <p>See also</p> <p><a href="#">includeHistory</a> [p 674]</p> <p><a href="#">includeDetails</a> [p 674]</p>                                                                                                                                                                                                                                                                                                             |
| includeMergeInfo    | <p>Includes the merge_info corresponding to a field of the technical data of an history transaction and has a potentially high access cost.</p> <p>Boolean type, default value is true.</p> <p>Note</p> <p>This parameter is ignored with the data category.</p> <p>See also <a href="#">REST access to history table</a> [p 301]</p>                                                                                   |
| includeSelector     | <p>Includes the selector field in the response, for each indirect reachable resource. The returned value corresponds to its URL resource.</p> <p>Type Boolean, default value is true.</p> <p>See also <a href="#">selector</a> [p 678]</p>                                                                                                                                                                              |
| includeSortCriteria | <p>Includes the sortCriteria field corresponding to the list of sort criteria applied. The sort criteria parameters are added by using:</p> <ul style="list-style-type: none"> <li>• <a href="#">sort</a> [p 677]</li> <li>• <a href="#">sortOnLabel</a> [p 677]</li> <li>• <a href="#">viewPublication</a> [p 678]</li> </ul> <p>Boolean type, default value is false.</p> <p>Example <a href="#">JSON</a> [p 709]</p> |
| includeTechnicals   | <p>Includes the internal technical data.</p> <p>Boolean type, default value is false.</p> <p>See also <a href="#">Technical data</a> [p 718]</p> <p>Note</p> <p>This parameter is ignored with the history category.</p>                                                                                                                                                                                                |
| includeValidation   | <p>Includes the validation report corresponding to the selection.</p> <p>Boolean type, default value is false.</p> <p>Note</p> <p>This parameter is ignored with the history category.</p> <p>See also <a href="#">includeDetails</a> [p 674]</p>                                                                                                                                                                       |

## Table parameters

The following parameters are applicable to tables, associations and selection nodes.

| Parameter             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| filter                | <p><a href="#">XPath predicate</a> [p 225] expression defines the field values to which the request is applied. If empty, all records will be retrieved.</p> <p>String type value.</p> <div> <p><b>Note</b></p> <p>The history code operation value is usable with ebx-operationCode path field from the meta section associated with this field.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| historyMode           | <p>Specifies the filter context applied on table.</p> <p>String type, possible values are:</p> <ul style="list-style-type: none"> <li>CurrentDataSpaceOnly: history in current dataspace</li> <li>CurrentDataSpaceAndAncestors: history in current dataspace and ancestors</li> <li>CurrentDataSpaceAndMergedChildren: history in current dataspace and merged children</li> <li>AllDataSpaces: history in all dataspace</li> </ul> <p>The default value is CurrentDataSpaceOnly.</p> <p>See also <a href="#">history</a> [p 661]</p> <div> <p><b>Note</b></p> <p>This parameter is ignored with the data category.</p> </div>                                                                                                                                                                                                                                                                                                |
| includeOcculting      | <p>Includes the records in occulting mode.</p> <p>Boolean type, default value is false.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| primaryKey            | <p>Search a record by a primary key, using the XPath expression. The <a href="#">XPath predicate</a> [p 225] expression should only contain field(s) of the primary key and all of them. Fields are separated by the operator and. A field is represented by one of the following possibilities according to its simple type:</p> <ul style="list-style-type: none"> <li>For the date, time or dateTime types: use the date-equal(path, value)</li> <li>For other types: indicate the path, the = operator and the value.</li> </ul> <p>Example with a composed primary key: ./pk1=1 and date-equal(./pk2d, '2015-11-13')</p> <p>The response will only contain the corresponding record, otherwise an error is returned. Consequently, the other table parameters are ignored (as <a href="#">filter</a> [p 676], <a href="#">viewPublication</a> [p 678], <a href="#">sort</a> [p 677], etc.)</p> <p>String type value.</p> |
| pageFirstRecordFilter | <p><i>Deprecated since version 5.9.0, replaced by pageRecordFilter</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| pageRecordFilter      | <p>Specifies the record <a href="#">XPath predicate</a> [p 225] expression filter of the page.</p> <p>String type value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pageAction  | <p>Specifies the pagination action to perform from page defined by pageRecordFilter.</p> <p>String type, possible values are:</p> <ul style="list-style-type: none"> <li>• first</li> <li>• previous</li> <li>• next</li> <li>• last</li> </ul> <p>See also <i>RequestPagination<sup>API</sup></i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| pageSize    | <p>Specifies the number of records per page.</p> <p>Integer type, default value is based on the user preferences.</p> <p>String type, the unbounded value can be defined to return all records. Only for this case, no pagination context is returned.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| sort        | <p>Specifies that the operation result will be sorted according to the specified criteria. The criteria are composed of one or more criteria, the result will be sorted by priority from the left. A criterion is composed of the field path and, optionally, the sorting order (ascending or descending, on value or on label). This parameter can be combined with:</p> <ol style="list-style-type: none"> <li>1. the <a href="#">sortOnLabel</a> [p 677] parameter as a new criteria added after the sort.</li> <li>2. the <a href="#">viewPublication</a> [p 678] parameter as a new criteria added after the sort.</li> </ol> <p>The value structure is as follows:<br/> &lt;path1&gt;:&lt;order&gt;;...;&lt;pathN&gt;:&lt;order&gt;</p> <p>Where:</p> <ul style="list-style-type: none"> <li>• &lt;path1&gt; corresponds to the field path in priority 1.</li> <li>• &lt;order&gt; corresponds to the sorting order, with one of the following values: <ul style="list-style-type: none"> <li>• asc: ascending order on value (default),</li> <li>• desc: descending order on value,</li> <li>• lasc: ascending order on label,</li> <li>• ldesc: descending order on label.</li> </ul> </li> </ul> <p>String type, the default value orders according to the primary key fields (ascending order on value).</p> <div> <p><b>Note</b></p> <p>The history code operation value is usable with the ebx-operationCode path field from the meta section associated with this field.</p> </div> <p>See also <i>Request.setSortCriteria<sup>API</sup></i></p> |
| sortOnLabel | <p>Specifies that the operation result will be sorted according to the record label. This parameter can be combined with:</p> <ol style="list-style-type: none"> <li>1. the <a href="#">sort</a> [p 677] parameter as a new criteria added before the sortOnLabel.</li> <li>2. the <a href="#">viewPublication</a> [p 678] parameter as a new criteria added after the sortOnLabel.</li> </ol> <p>The value structure is as follows:<br/> &lt;order&gt;</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

| Parameter       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                 | <p>Where:</p> <ul style="list-style-type: none"> <li>• &lt;order&gt; corresponds to the sorting order, with one of the following values: <ul style="list-style-type: none"> <li>• 1asc: ascending order on label,</li> <li>• 1desc: descending order on label.</li> </ul> </li> </ul> <p>The behavior of this parameter is described in the section <a href="#">defaultLabel</a> [p 497].</p> <p>String type value.</p> <p>See also <a href="#">Limitation</a> [p 694]</p>                                                                                                                    |
| viewPublication | <p>Specifies the name of the published view. This parameter can be combined with:</p> <ol style="list-style-type: none"> <li>1. the <a href="#">filter</a> [p 676] parameter as the logical and operation.</li> <li>2. the <a href="#">sort</a> [p 677] parameter as a new criteria added before the viewPublication.</li> <li>3. the <a href="#">sortOnLabel</a> [p 677] parameter as new criteria added before the viewPublication.</li> </ol> <p>The behavior of this parameter is described in the section <a href="#">EBX® as a Web Component</a> [p 194].</p> <p>String type value.</p> |

## Selector parameters

The following parameters are only applicable to fields that return an enumeration, foreign key or osd:resource (Example [JSON](#) [p 716]). By default, a pagination mechanism is enabled.

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| selector          | <p>Specifies whether:</p> <ul style="list-style-type: none"> <li>• true: returns all possible values (includes their labels)</li> <li>• false: returns the current values for the current field.</li> </ul> <p>Boolean type, default value is false.</p> <div> <p><b>Note</b></p> <p>This parameter is ignored with the history category.</p> </div> |
| firstElementIndex | <p>Specifies the index of the first element returned by the selector. Must be an integer higher than or equal to 0.</p> <p>Integer type, default value is 0.</p>                                                                                                                                                                                     |
| pageSize          | <p>Specifies the number of elements per page.</p> <p>Integer type, default value is based on the user preferences.</p> <p>String type, the unbounded value can be defined to return all values. Only for this case, no pagination context is returned.</p>                                                                                           |
| selectorFilter    | <p>Specifies the filter of the selector.</p> <p>String type value, the syntax complies with the <a href="#">Text search</a> [p 112].</p>                                                                                                                                                                                                             |

## HTTP codes

| HTTP code                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 ( <i>OK</i> )          | The selected resource is successfully retrieved.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 400 ( <i>Bad request</i> ) | <p>The request is incorrect. This occurs when:</p> <ul style="list-style-type: none"> <li>• The selected field in a record or a dataset is sub-terminal,</li> <li>• The XPath predicate of the <code>filter</code> parameter is malformed or contains unfilterable nodes.</li> <li>• The XPath predicate of the <code>primaryKey</code> parameter is malformed or is not a record primary key.</li> <li>• The sort criteria of the <code>sort</code> parameter have an invalid syntax or contain unsortable nodes.</li> <li>• <code>pageAction</code> parameter value is not included in allowed values, or <code>pageRecordFilter</code> is malformed or non-existent when selecting next or previous page.</li> <li>• <code>pageSize</code> parameter value is below 2, or is a string different from unbounded.</li> <li>• The table view for the <code>viewPublication</code> parameter is either hierarchical, non-existent or non-published.</li> <li>• The <code>selector</code> parameter is used for a non-enumerated node, or the <code>firstElementIndex</code> is negative, higher than or equal to the number of values.</li> </ul> |
| 403 ( <i>Forbidden</i> )   | The selected resource is hidden for the authenticated user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 404 ( <i>Not found</i> )   | The selected resource is not found.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Response body

After a successful dataset, table, record or field selection, the result is returned in the response body. The content depends on the provided parameters and selected data.

Example: [JSON](#) [p 701].

## Insert operation

Insert operation uses the `POST` HTTP method. A body message is required to specify data. This operation supports the insertion of one or more records in a single transaction. Moreover, it is also possible to update record(s) through parameterization.

- **Record:** insert a new record or modify an existing one in the selected table.
- **Record table:** insert or modify one or more records in the selected table, while securing a consistent answer. Operations are executed sequentially, in the order defined on the client side. When an error occurs during a table operation, all updates are cancelled and the client receives an error message with detailed information.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/{dataset}/{pathInDataset}
```

Where:

- `{dataspace}` corresponds to `B` followed by the dataspace identifier or to `V` followed by the snapshot identifier.

- {dataset} corresponds to the dataset identifier.
- {pathInDataset} corresponds to the path of the table node.



## Parameters

The following parameters are applicable with the insert operation.

| Parameter                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>includeDetails</code>              | <p>Includes the <code>details</code> field in the answer for access to the details of the data. The returned value corresponds to its URL resources.</p> <p>Type <code>Boolean</code>, the default value is <code>false</code>.</p> <div> <p><b>Note</b></p> <p>Only applicable on the <b>record table</b>.</p> </div>                                                                                                                                                                                                                                       |
| <code>includeForeignKey</code>           | <p>Includes the <code>foreignKey</code> field in the answer for each record. The returned value corresponds to the value of a foreign key field that was referencing this record.</p> <p><code>Boolean</code> type, the default value is <code>false</code>.</p> <div> <p><b>Note</b></p> <p>Only applicable on the <b>record table</b>.</p> </div>                                                                                                                                                                                                          |
| <code>includeLabel</code>                | <p>Includes the <code>label</code> field in the answer for each record.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> <li>• <code>yes</code>: the <code>label</code> field is included.</li> <li>• <code>no</code>: the <code>label</code> field is not included (use case: integration).</li> </ul> <p><code>String</code> type, the default value is <code>no</code>.</p> <div> <p><b>Note</b></p> <p>Only applicable on the <b>record table</b>.</p> </div>                                                                          |
| <code>updateOrInsert</code>              | <p>Specifies the behavior when the record to insert already exists:</p> <ul style="list-style-type: none"> <li>• If <code>true</code>: the existing record is updated with new data.</li> </ul> <p>For a request on a <b>record table</b>, the <code>code</code> field is added to the report in order to specify if this is an insert 201 or an update 204.</p> <ul style="list-style-type: none"> <li>• If <code>false</code> (default value): a client error is returned and the operation is aborted.</li> </ul> <p><code>Boolean</code> type value.</p> |
| <code>byDelta</code>                     | <p>Specifies the behavior for setting value of nodes that are not defined in the request body. This is described in the <a href="#">Update modes</a> [p 718] section.</p> <p><code>Boolean</code> type, the default value is <code>true</code>.</p> <div> <p><b>Note</b></p> <p>Applicable on record in update mode and if the <a href="#">updateOrInsert</a> [p 681] parameter is <code>true</code>.</p> </div>                                                                                                                                             |
| <code>blockingConstraintsDisabled</code> | <p>Specifies whether blocking constraints are ignored, if so, the operation is committed regardless of the validation error created, otherwise, the operation would be aborted.</p> <p><code>Boolean</code> type, default value is <code>false</code>.</p>                                                                                                                                                                                                                                                                                                   |

| Parameter | Description                                                                             |
|-----------|-----------------------------------------------------------------------------------------|
|           | See <a href="#">Blocking and non-blocking constraints</a> [p 524] for more information. |

## Message body

The request must define a message body. The format depends on the inserted object type:

- **Record:** similar to the select operation of a record but without the record's header (example [JSON](#) [p 697]).
- **Record table:** Similar to the select operation on a table but without the pagination information (example [JSON](#) [p 698]).

See also [Inheritance](#) [p 692]

## HTTP codes

| HTTP code                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 (OK)                   | If the request relates to a <b>record table</b> .<br>The insert request was applied successfully, an optional report is returned in the response body.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 201 (Created)              | If the request relates to a <b>record</b> .<br>A new record has been created, in this case, the header field <code>Location</code> is returned with its resource URL.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 204 (No content)           | If the request relates to a <b>record</b> .<br>Only available if <code>updateOrInsert</code> is true, an existing record has been successfully updated, in this case, the header field <code>Location</code> is returned with its resource URL.                                                                                                                                                                                                                                                                                                                                                 |
| 400 (Bad request)          | The request is incorrect. This occurs when the body message structure does not comply with what was mentioned in <a href="#">Message body</a> [p 682].                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 403 (Forbidden)            | Authenticated user is not allowed to create a record or the request body contains a read-only field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 404 (Not found)            | The selected resource is not found.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 409 (Conflict)             | Concurrent modification, only available if <code>updateOrInsert</code> is true, the <a href="#">Optimistic locking</a> [p 691] is activated and the content has changed in the meantime, it must be reloaded before update.                                                                                                                                                                                                                                                                                                                                                                     |
| 422 (Unprocessable entity) | The request cannot be processed. This occurs when: <ul style="list-style-type: none"> <li>A blocking validation error occurs (only available if <code>blockingConstraintsDisabled</code> is false).</li> <li>The record cannot be inserted because a record with the same primary key already exists (only available if <code>updateOrInsert</code> is false).</li> <li>The record cannot be inserted because the definition of the primary key is either non-existent or incomplete.</li> <li>The record cannot be updated because the value of the primary key cannot be modified.</li> </ul> |

## Response body

The response body format depends on the inserted object type:

- **Record:** is empty if the operation was executed successfully. The header field `Location` is returned with its URL resource.
- **Record table:** (optional) contains a table of element(s), corresponding to the insert operation report (example [JSON](#) [p 716]). This report is automatically included in the response body, if at least one of the following options is set:
  - `includeForeignKey`
  - `includeLabel`
  - `includeDetails`

See also [Inheritance](#) [p 692]

## Update operation

This operation allows the modification of a single dataset or record. The PUT HTTP method must be used. Available URL formats are:

- **Dataset node:** modifies the values of terminal nodes contained in the selected node.  
`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/{dataset}/{pathInDataset}`
- **Record:** modifies the content of selected record.

### Note

Also available for POST HTTP methods. In this case, the URL must correspond to the table by setting the parameter `updateOrInsert` to true.

- **Field:** modifies the field content.  
`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}/{pathInRecord}`

### Note

The field must be either a terminal node or above.

Where:

- {dataspace} corresponds to B followed by the dataspace identifier or to v followed by the snapshot identifier.
- {dataset} corresponds to the dataset identifier.
- {pathInDataset} corresponds to the path of the dataset node:
  - For dataset node operations, this must be any terminal node or above except table node,
  - For record and field operations, this corresponds to the table node.
- {encodedPrimaryKey} corresponds to the encoded representation of the primary key.

See also *RESTEncodingHelper*<sup>API</sup>

- {pathInRecord} corresponds to the path starting from the table node.

## Parameters

Here are the parameters applicable with the update operation.

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| blockingConstraintsDisabled        | Specifies whether blocking constraints are ignored, if so, the operation is committed regardless of the validation error created, otherwise, the operation would be aborted.<br>Boolean type, default value is <code>false</code> .<br>See <a href="#">Blocking and non-blocking constraints</a> [p 524] for more information. |
| byDelta                            | Specifies the behavior for setting value of nodes that are not defined in the request body. This is described in the <a href="#">Update modes</a> [p 718] section.<br>Boolean type, the default value is <code>true</code> .                                                                                                   |
| checkNotChangedSinceLastUpdateDate | Timestamp in datetime format used to ensure that the record has not been modified since the last read. Also see the <a href="#">Optimistic locking</a> [p 691] section.<br>DateTime type value.                                                                                                                                |

## Message body

The request must define a message body.

The structure is the same as for:

- the dataset node (sample [JSON](#) [p 696]),
- the record (sample [JSON](#) [p 697]),
- the record fields (sample [JSON](#) [p 697]),

depending on the updated scope, by only keeping the content entry.

See also [Inheritance](#) [p 692]

## HTTP codes

| HTTP code                           | Description                                                                                                                                                                                                                                                                                                                |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 204 ( <i>No content</i> )           | The record, field or dataset node has been successfully updated.                                                                                                                                                                                                                                                           |
| 400 ( <i>Bad request</i> )          | The request is incorrect. This occurs when the body request structure does not comply.                                                                                                                                                                                                                                     |
| 403 ( <i>Forbidden</i> )            | Authenticated user is not allowed to update the specified resource or the request body contains a read-only field.                                                                                                                                                                                                         |
| 404 ( <i>Not found</i> )            | The selected resource is not found.                                                                                                                                                                                                                                                                                        |
| 409 ( <i>Conflict</i> )             | Concurrent modification, the <a href="#">Optimistic locking</a> [p 691] is activated and the content has changed in the meantime, it must be reloaded before the update.                                                                                                                                                   |
| 422 ( <i>Unprocessable entity</i> ) | The request cannot be processed. This occurs when: <ul style="list-style-type: none"> <li>A blocking validation error occurs (only available if <code>blockingConstraintsDisabled</code> is <code>false</code>).</li> <li>The record cannot be updated because the value of the primary key cannot be modified.</li> </ul> |

## Delete operation

The operation uses the DELETE HTTP method.

Two URL formats are available:

- **Record:** delete a record specified in the URL.  
`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}`
- **Record table:** deletes several records in the specified table, while providing a consistent answer. This mode requires a body message containing a record table. The deletions are executed sequentially, according to the order defined in the table. When an error occurs during a table operation, all deletions are cancelled and an error message is displayed with detailed information.  
`http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/{dataset}/{pathInDataset}`

Where:

- {dataspace} corresponds to B followed by the dataspace identifier or to v followed by the snapshot identifier.
- {dataset} corresponds to the dataset identifier.
- {pathInDataset} corresponds to the path of the table node.
- {encodedPrimaryKey} corresponds to the encoded representation of the primary key.

See also `RESEncodingHelper`<sup>API</sup>

In a child dataset context, this operation modifies the `inheritanceMode` property value of the record as follows:

- A record with inheritance mode set to `inherit` or `overwrite` becomes `occult`.
- A record with inheritance mode set to `occult` becomes `inherit` if the `inheritIfInOccultingMode` operation parameter is set to `true` or is undefined, otherwise there is no change.
- A record with inheritance mode set to `root` is simply deleted.

See also [Inheritance](#) [p 692]

## Parameters

Here are the following parameters applicable with delete operation.

| Parameter                                       | Description                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>includeOcculting</code>                   | Includes occulted records.<br>Boolean type, the default value is <code>false</code> .                                                                                                                                                                                                                                          |
| <code>inheritIfInOccultingMode</code>           | <i>Deprecated since version 5.8.1</i> While it remains available for backward compatibility reasons, it will eventually be removed in a future version.<br>Inherits the record if it is in occulting mode.<br>Boolean type, the default value is <code>true</code> .                                                           |
| <code>checkNotChangedSinceLastUpdateDate</code> | Timestamp in datetime format used to ensure that the record has not been modified since the last read. Also see the <a href="#">Optimistic locking</a> [p 691] section.<br>DateTime type value.                                                                                                                                |
| <code>blockingConstraintsDisabled</code>        | Specifies whether blocking constraints are ignored, if so, the operation is committed regardless of the validation error created, otherwise, the operation would be aborted.<br>Boolean type, default value is <code>false</code> .<br>See <a href="#">Blocking and non-blocking constraints</a> [p 524] for more information. |

## Message body

The request must define a message body only when deleting several records:

- **Record table:** The message contains a table of elements related to a record, with for each element one of the following properties:
  - `details`: corresponds to the record URL, it is returned by the select operation.
  - `primaryKey`: corresponds to the primary key of the record, using the XPath expression.
  - `foreignKey`: corresponds to the value that a foreign key would have if it referred to a record.

See also `PrimaryKey`<sup>API</sup>

Example [JSON](#) [p 698].

## HTTP codes

| HTTP code                           | Description                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 ( <i>OK</i> )                   | The operation has been executed successfully. A report is returned in the response body.                                                                                                                                                                                                                                              |
| 400 ( <i>Bad request</i> )          | The request is incorrect. This occurs when: <ul style="list-style-type: none"> <li>the structure of the message body does not comply with <a href="#">Message body</a> [p 687].</li> <li>the message body contains a record table while the URL specifies a record.</li> </ul>                                                        |
| 403 ( <i>Forbidden</i> )            | Authenticated user is not allowed to delete or occult the specified record.                                                                                                                                                                                                                                                           |
| 404 ( <i>Not found</i> )            | The selected record is not found. In the child dataset context, it should be necessary to use the <code>includeOcculting</code> parameter.                                                                                                                                                                                            |
| 409 ( <i>Conflict</i> )             | Concurrent modification, The <a href="#">Optimistic locking</a> [p 691] is activated and the content has changed in the meantime, it must be reloaded before deleting the record.<br>The parameter value <code>checkNotChangedSinceLastUpdateDate</code> exists but does not correspond to the actual last update date of the record. |
| 422 ( <i>Unprocessable entity</i> ) | Only available if <code>blockingConstraintsDisabled</code> is <code>false</code> , the operation fails because of a blocking validation error.                                                                                                                                                                                        |

## Response body

After a successful record deletion or occulting, a report is returned in the response body. It contains the number of deleted, occulted and inherited record(s).

Example [JSON](#) [p 717].

## Count operation

Count operation may use one of the following methods:

- GET HTTP method,
- POST HTTP method without message body or
- POST HTTP method with message body but without content field on root.

The URL formats are:

- Dataset node:** the data category returns the number of terminal nodes contained in the selected node.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/{dataset}/{pathInDataset}?count=true
```

### Note

Not applicable with the history category.

- Table** depending on the [operation category](#) [p 661]:  
the data category returns the number of table records.



The history category returns the number of table history records.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/{dataspace}/{dataset}/{pathInDataset}?count=true
```

- **Field** depending on the [operation category](#) [p 661]:

the data category counts the record fields.

The history category counts the history record field.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}/{pathInRecord}?count=true
```

#### Note

The field must be either an association node, a selection node, a terminal node or above.

Where:

- {category} corresponds to the [operation category](#) [p 661].
- {dataspace} corresponds to B followed by the dataspace identifier or to v followed by the snapshot identifier.
- {dataset} corresponds to the dataset identifier.
- {pathInDataset} corresponds to the path of the dataset node, that can be a group node or a table node.
- {encodedPrimaryKey} corresponds to the encoded representation of the primary key.

See also *RESEncodingHelper*<sup>API</sup>

- {pathInRecord} corresponds to the path starting from the table node.

## Parameters

The following parameters are applicable to the count operation.

| Parameter | Description                                                                                                                            |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------|
| count     | This parameter is used to specify whether this is a count operation or a selection operation.<br>Boolean type, default value is false. |

## Table parameters

The following parameters are applicable to tables, associations and selection nodes.

| Parameter        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| filter           | <p><a href="#">XPath predicate</a> [p 225] expression defines the field values to which the request is applied. If empty, all records will be considered.</p> <p>String type value.</p> <div> <p>Note</p> <p>The history code operation value is usable with the ebx-operationCode path field from the meta section associated with this field.</p> </div>                                                                                                                                                                                                                                                              |
| historyMode      | <p>Specifies the filter context applied on table.</p> <p>String type, possible values are:</p> <ul style="list-style-type: none"> <li>CurrentDataSpaceOnly: history in current dataspace</li> <li>CurrentDataSpaceAndAncestors: history in current dataspace and ancestors</li> <li>CurrentDataSpaceAndMergedChildren: history in current dataspace and merged children</li> <li>AllDataSpaces: history in all dataspace</li> </ul> <p>The default value is CurrentDataSpaceOnly.</p> <p>See also <a href="#">history</a> [p 661]</p> <div> <p>Note</p> <p>This parameter is ignored with the data category.</p> </div> |
| includeOcculting | <p>Includes the records in occulting mode.</p> <p>Boolean type, default value is false.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| viewPublication  | <p>Specifies the name of the published view to be considered during the count execution. This parameter can be combined with:</p> <ul style="list-style-type: none"> <li>the <a href="#">filter</a> [p 690] parameter as the logical and operation.</li> </ul> <p>The behavior of this parameter is described in the section <a href="#">EBX® as a Web Component</a> [p 194].</p>                                                                                                                                                                                                                                       |

## Selector parameters

The following parameters are only applicable to fields that return an enumeration, foreign key or `osd:resource`.

| Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>selector</code>       | <p>Specifies whether:</p> <ul style="list-style-type: none"> <li><code>true</code>: returns the number of all possible values</li> <li><code>false</code>: returns the number of possible values for the current field.</li> </ul> <p>Boolean type, default value is <code>false</code>.</p> <div> <p><b>Note</b></p> <p>This parameter is ignored with the history category.</p> </div> |
| <code>selectorFilter</code> | <p>Specifies the filter of the selector.</p> <p>String type value, the syntax complies with the <a href="#">Text search</a> [p 112].</p>                                                                                                                                                                                                                                                 |

## HTTP codes

| HTTP code                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 200 ( <i>OK</i> )          | The selected resource is successfully counted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 400 ( <i>Bad request</i> ) | <p>The request is incorrect. This occurs when:</p> <ul style="list-style-type: none"> <li>The selected field in a record or a dataset is sub-terminal.</li> <li>The selected dataset field is a dataset tree.</li> <li>The XPath predicate of the <code>filter</code> parameter is malformed or contains unfilterable nodes.</li> <li>The table view for the <code>viewPublication</code> parameter is either hierarchical, non-existent or non-published.</li> <li>The <code>selector</code> parameter is used for a non-enumerated node, or the <code>firstElementIndex</code> is negative, higher than or equal to the number of values.</li> </ul> |
| 403 ( <i>Forbidden</i> )   | The selected resource is hidden for the authenticated user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 404 ( <i>Not found</i> )   | The selected resource is not found.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Optimistic locking

To prevent an update or a delete operation on a record that was previously read but may have changed in the meantime, an optimistic locking mechanism is provided.

To enable optimistic locking, a select request must set the parameter `includeTechnicals` to `true`.

See [Technical data](#) [p 718] for more information.

The value of the `lastUpdateDate` property must be included in the following update request. If the record has been changed since the specified time, the update or delete will be cancelled.

- **Record:** update whole or partial content of the selected record.

The property `lastUpdateDate` should be added to the request body to prevent update of a modified record.

See the [JSON](#) [p 697] example of a record.

- **Field:** update of a single field of the selected record.

The property value `lastUpdateDate` must be declared in the request URL by the `checkNotChangedSinceLastUpdateDate` parameter to prevent the update of a modified record.

The property value `lastUpdateDate` can also be used in the request URL `checkNotChangedSinceLastUpdateDate` parameter to prevent deletion on a modified record.

#### Note


The `checkNotChangedSinceLastUpdateDate` parameter may be used more than once but only on the same record. This implies that if the request URL returns more than one record, the request will fail.

## Inheritance

EBX® inheritance features are supported by built-in RESTful services using specific properties and automatic behaviors. In most cases, the inheritance state will be automatically computed by the server according to the record and field definition or content. Every action that modifies a record or a field may have an indirect impact on those states. In order to fully handle the inheritance life cycle, direct modifications of the state are allowed under certain conditions. Forbidden or incoherent explicit alteration attempts are ignored.

See also [Inheritance and value resolution](#) [p 268]

### Record inheritance life cycle in built-in RESTful services



## Inheritance properties

The following table describes properties related to the EBX® inheritance features.

| Property           | Location                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| inheritance        | record or table metadata  | <p>Specifies if dataset inheritance is activated for the table. The value is computed from the data model and cannot be modified through built-in RESTful services.</p> <p>See also <a href="#">inheritance property in metadata</a>. [p 707]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| inheritedField     | field metadata            | <p>Specifies the field's value source. The source data are directly taken from the data model and cannot be modified through built-in RESTful services.</p> <p>See also <a href="#">inheritedField property in metadata</a>. [p 708]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| inheritanceMode    | record in child dataset   | <p>Specifies the record's inheritance state. To set a record's inheritance from overwrite to inherit, its inheritanceMode value must be explicitly provided in the request. In this specific case, the content property will be ignored if present. occult and root explicit values are always ignored. An overwrite explicit value without a content property is ignored.</p> <div> <p><b>Note</b></p> <p>Inherited record's fields are necessarily inherit.</p> </div> <div> <p><b>Note</b></p> <p>Root records in child dataset will always be root.</p> </div> <p>Possible values are: root, inherit, overwrite, occult. For more information, see <a href="#">Record lookup mechanism</a> [p 270].</p> |
|                    | field in overwrite record | <p>Specifies the field's inheritance state. To set a field's inheritance to inherit, its inheritanceMode value must be explicitly provided in the request. The content property will be ignored in this case. overwrite explicit value without a content property is ignored.</p> <div> <p><b>Note</b></p> <p>inheritanceMode at field level does not appear for root, inherit and occult records.</p> </div> <div> <p><b>Note</b></p> <p>inheritedFieldMode and inheritanceMode properties cannot be both set on the same field.</p> </div> <p>Possible values are: inherit, overwrite. For more information, see <a href="#">Inheritance and value resolution</a> [p 268].</p>                            |
| inheritedFieldMode | inherited field           | <p>Specifies the inherited field's inheritance state. To set a field's inheritance to inherit, its inheritedFieldMode value must be</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Property | Location | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          |          | <p>explicitly provided in the request. The content property will be ignored in this case. overwrite explicit values without a content property are ignored.</p> <p><b>Note</b></p> <p>inheritedFieldMode and inheritanceMode properties cannot be both set on the same field.</p> <p><b>Note</b></p> <p>inheritedFieldMode has priority over the inheritanceMode property.</p> <p>Possible values are: inherit, overwrite. For more information, see <a href="#">Value lookup mechanism</a> [p 271].</p> |

## 99.7 Limitations

### **General limitations**

- Indexes, in the request URL {pathInDataset} or {pathInRecord}, are not supported.
- Nested aggregated lists are not supported.
- Dataset nodes and field operations applied to nodes that are sub-terminal are not supported.

See [Access properties](#) [p 539] for more information about terminal nodes.

### **Read operations**

- Within the selector, the pagination context is limited to the nextPage property.
- Within the viewPublication parameter, the hierarchical view is not supported.
- The sortOnLabel parameter ignores programmatic labels.
- The system information response's properties cannot be browsed through the REST URL with the hierarchical representation.

See [System information operation](#) [p 667] for more information.

### **Write operations**

- Association fields cannot be updated, therefore, the list of associated records cannot be modified directly.
- Control policy onUserSubmit-checkModifiedValues of the user interface is not supported. To retrieve validation errors, invoke the select operation on the resource by including the includeValidation parameter.

See [Blocking and non-blocking constraints](#) [p 524] for more information.

### **Directory operations**

- Changing or resetting a password for a user is not supported.

# CHAPTER 100

---

## JSON format

This chapter contains the following topics:

1. [Introduction](#)
2. [Global structure](#)
3. [Meta-data](#)
4. [Sort criteria information](#)
5. [Validation](#)
6. [Content](#)
7. [Update modes](#)
8. [Known limitations](#)

### 100.1 Introduction

The JSON (JavaScript Object Notation) is the data-interchange format used by TIBCO EBX® [RESTful operations](#) [p 659].

This format is lightweight, self-describing and can be used to design UIs or to integrate EBX® in a company's information system.

- The data context is exhaustive, except for association fields and selection nodes that are not directly returned. However, these fields are included in the response with a URL link named `details` included by default. It can be indirectly used to get the fields content.
- The volume of data is limited by a pagination mechanism activated by default, it can be configured or disabled.

URL formatted links allow retrieving:

- Tables, records, dataset non-terminal nodes, foreign keys, resource fields (property details).
- Possible values for foreign keys or enumerations (selector parameter).

See also [Activation and configuration](#) [p 652]

#### Note

JSON data are always encoded in UTF-8.

## 100.2 Global structure

### ***JSON Request body***

Request body is represented by a JSON object whose content varies according to the operation and its category.

#### **Auth category**

The request body holds several properties directly placed in the root JSON object.

- **Token creation**

Specifies the login and password to use for an authentication token creation attempt.

```
{
  "login": "...",                // JSON String
  "password": "...",            // JSON String
}
```

Specifies the specific attribute, to activate the user authentication against the HTTP request, for an authentication token creation attempt.

```
{
  "specific": true                // JSON Boolean
}
```

See also [Create token operation](#) [p 668]

- **Password change**

Specifies the login, password and passwordNew to use for the password change.

```
{
  "login": "...",                // JSON String
  "password": "...",            // JSON String
  "passwordNew": "...",          // JSON String
}
```

See also [Change password operation](#) [p 670]

#### **Data category**

The request body contains at least a content property where master data values will be defined.

- **Dataset node**

Specifies the target values of terminal nodes under the specified node. This request is used on the dataset node update operation.

```
{
  "content": {
    "nodeName1": {
      "content": true
    },
    "nodeName2": {
      "content": 2
    },
    "nodeName3": {
      "content": "Hello"
    }
  }
}
```



See also [Update operation](#) [p 684]

- **Record**

Specifies the target record content by setting the value for each field. For missing fields, the behavior depends on the request parameter `byDelta`. This structure is used on table record insert or on record update.

See also [Inheritance](#) [p 692]

Some technical data can be added beside the content property such as `lastUpdateDate`.

See also [Optimistic locking](#) [p 691]

```
{
  ...
  "lastUpdateDate": "2015-12-25T00:00:00.001",
  ...
  "content": {
    "gender": {
      "content": "Mr."
    },
    "lastName": {
      "content": "Chopin"
    },
    "lastName-en": {
      "content": "Chopin",
      "inheritedFieldMode": "inherit"
    },
    "firstName": {
      "content": "Fryderyk"
    },
    "firstName-en": {
      "content": "Frédéric",
      "inheritedFieldMode": "overwrite"
    },
    "birthDate": {
      "content": "1810-03-01"
    },
    "deathDate": {
      "content": "1849-10-17"
    },
    "jobs": {
      "content": [
        {
          "content": "CM"
        },
        {
          "content": "PI"
        }
      ]
    },
    "infos": {
      "content": [
        {
          "content": "https://en.wikipedia.org/wiki/Chopin"
        }
      ]
    }
  }
}
```

See also

[Insert operation](#) [p 679]

[Update operation](#) [p 684]

- **Record fields**

Specifies the target values of fields under the record terminal node by setting the value of each field. For missing fields, the behavior depends on the request parameter `byDelta`. This structure is only used for table record updates.

See also [Inheritance](#) [p 692]

```
{
  "content": [
    {
      "content": "CM"
    },
    {
      "content": "PI"
    }
  ]
}
```

See also [Update operation](#) [p 684]

- **Record table**

Defines the content of one or more records by indicating the value of each field. For missing fields, the behavior depends on the parameter of the request `byDelta`. This structure is used upon insert or update of records in the table.

```
{
  "rows": [
    {
      "content": {
        "gender": {
          "content": "M"
        },
        "lastName": {
          "content": "Saint-Saëns"
        },
        "firstName": {
          "content": "Camille"
        },
        "birthDate": {
          "content": "1835-10-09"
        },
        ...
      }
    },
    {
      "content": {
        "gender": {
          "content": "M"
        },
        "lastName": {
          "content": "Debussy"
        },
        "firstName": {
          "content": "Claude"
        },
        "birthDate": {
          "content": "1862-10-22"
        },
        ...
      }
    }
  ]
}
```

See also

[Insert operation](#) [p 679]

[Update operation](#) [p 684]

- **Record table to be deleted**

Defines one or more records. This structure is used upon deleting several records from the same table.

```
{
  "rows": [
    {
      "details": "http://.../root/table/1"
    },
    {
      "details": "http://.../root/table/2"
    },
    {
      "primaryKey": ".oid=3"
    },
    {
      "foreignKey": "4"
    },
    ...
  ]
}
```

See also [Delete operation](#) [p 686]

- **Field**

Specifies the target field content. This request is used on the field update.

The request has the same structure as defined in [node value](#) [p 711] by only keeping the content entry. Other entries are simply ignored.

See also [Update operation](#) [p 684]

- **Open or close user interface**

Specifies whether the user interface is open or close and the unavailability message.

```
{
  "content": {
    "toolStatus": {
      "content": true // or false
    },
    "toolStatusCloseMessage": {
      "content": "Access is temporarily forbidden for maintenance."
    }
  }
}
```

See also [User interface operations](#) [p 667]

Only writable fields can be mentioned in the request, this excludes the following cases:

- Association node,
- Selection node,
- Value function,
- JavaBean field that does not have a setter,
- Unwritable permission on node for authenticated user.

## **JSON Response body**

The response body is represented by a JSON object whose content depends on the operation and its category.

## Admin category

The selection operation for this category only provides the values corresponding to the request under a content property.

- **System information**

Contains EBX® instance's system information. The representation of these data can be flat or hierarchical.

Flat representation:

```
{
  "content": {
    "bootInfoEBX": {
      "label": "EBX® configuration",
      "content": {
        "product.version": {
          "label": "EBX® product version",
          "content": "5.8.1 [...] Enterprise Edition"
        },
        "product.configuration.file": {
          "label": "EBX® main configuration file",
          "content": "System property [ebx.properties=../ebx.properties]"
        }
      }
      // others keys
    },
    "repositoryInfo": {
      "label": "Repository information",
      "content": {
        "repository.identity": {
          "label": "Repository identity",
          "content": "00905A5753FD"
        },
        "repository.label": {
          "label": "Repository label",
          "content": "My repository"
        }
      }
      // others keys
    },
    "bootInfoVM": {
      "label": "System information",
      "content": {
        "java.home": {
          "label": "Java installation directory",
          "content": "C:\\JTools\\jdk1.8.0\\jre"
        },
        "java.vendor": {
          "label": "Java vendor",
          "content": "Oracle Corporation"
        }
      }
      // others keys
    }
  }
}
```

Hierarchical representation:

```
{
  "content": {
    "bootInfoEBX": {
      "label": "EBX® configuration",
      "content": {
        "product": {
          "content": {
            "version": {
              "label": "EBX® product version",
              "content": "5.8.1 [...] Enterprise Edition"
            }
          },
          "configuration": {
            "content": {
              "file": {
                "label": "EBX® main configuration file",
                "content": "System property [ebx.properties=../ebx.properties]"
              }
            }
          }
        }
      }
    }
  }
}
```

```

    },
    "vm": {
      "content": {
        "startTime": {
          "label": "VM start time",
          "content": "2017/09/11-10:04:17-0729 CEST"
        },
        "identifier": {
          "label": "VM identifier",
          "content": "1"
        }
      }
    },
    // other hierarchical keys
  }
}

```

See also [System information operation](#) [p 667]

## Auth category

The response body contains several properties directly placed in its root JSON object.

- **Token creation**

Contains the token value and its type.

```

{
  "accessToken": "...",           // JSON String
  "tokenType": "...",           // JSON String
}

```

See also [Create token operation](#) [p 668]

## Data category

The selection operation contains two different parts.

The first one named meta contains the exhaustive structure of the response.

The other, regrouping content, rows, pagination...etc, contains the values corresponding to the request.

- **Dataset tree**

Contains the hierarchy of table and non-terminal group nodes.

```

{
  "meta": {
    "fields": [
      {
        "name": "rootName",
        "label": "Localized label",
        "description": "Localized description",
        "type": "group",
        "pathInDataset": "/rootName",
        "fields": [
          {
            "name": "settings",
            "label": "Settings",
            "type": "group",
            "pathInDataset": "/rootName/settings",
            "fields": [
              {
                "name": "settingA",
                "label": "A settings label",
                "type": "group",
                "pathInDataset": "/rootName/settings/settingA"
              },
              {
                "name": "settingB",

```

```

        "label": "B settings label",
        "type": "group",
        "pathInDataset": "/rootName/settings/settingB"
    }
}
],
{
    "name": "table1",
    "label": "Table1 localized label",
    "type": "table",
    "minOccurs": 0,
    "maxOccurs": "unbounded",
    "pathInDataset": "/rootName/table1"
},
{
    "name": "table2",
    "label": "Table2 localized label",
    "type": "table",
    "minOccurs": 0,
    "maxOccurs": "unbounded",
    "pathInDataset": "/rootName/table2"
}
]
}
}],
"validation": [
    {
        "level": "error",
        "message": "Value must be greater than or equal to 0.",
        "details": "http://.../rootName/settings/settingA/settingA1?includeValidation=true"
    },
    {
        "level": "error",
        "message": "Field 'Settings A2' is mandatory.",
        "details": "http://.../rootName/settings/settingA/settingA2?includeValidation=true"
    }
],
"content": {
    "rootName": {
        "details": "http://.../rootName",
        "content": {
            "settings": {
                "details": "http://.../rootName/settings",
                "content": {
                    "weekTimeSheet": {
                        "details": "http://.../rootName/settings/settingA"
                    },
                    "vacationRequest": {
                        "details": "http://.../rootName/settings/settingB"
                    }
                }
            },
            "table1": {
                "details": "http://.../rootName/table1"
            },
            "table2": {
                "details": "http://.../rootName/table2"
            }
        }
    }
}
}
}

```

The meta and validation properties are optional.

## See also

*Meta-data* [p 706]

## Validation [p 710]

Select operation [p 671]

- **Dataset node**

Contains the list of terminal nodes under the specified node.

```
{
  "meta": {
    "fields": [
      {
        "name": "id",
        "type": "string",
        "required": true,
        "indexed": true,
        "unique": true
      },
      {
        "name": "name",
        "type": "string",
        "required": true,
        "indexed": true,
        "unique": false
      },
      {
        "name": "age",
        "type": "integer",
        "required": false,
        "indexed": false,
        "unique": false
      },
      {
        "name": "email",
        "type": "string",
        "required": false,
        "indexed": true,
        "unique": false
      },
      {
        "name": "password",
        "type": "string",
        "required": true,
        "indexed": false,
        "unique": true
      },
      {
        "name": "created_at",
        "type": "timestamp",
        "required": true,
        "indexed": true,
        "unique": false
      },
      {
        "name": "updated_at",
        "type": "timestamp",
        "required": true,
        "indexed": true,
        "unique": false
      }
    ]
  }
}
```

```

        "name": "nodeName1",
        "label": "Localized label of the field node 1",
        "description": "Localized description",
        "type": "boolean",
        "minOccurs": 1,
        "maxOccurs": 1,
        "pathInDataset": "/rootName/.../nodeName1"
      },
      {
        "name": "nodeName2",
        "label": "Localized label of the field node 2",
        "type": "int",
        "minOccurs": 1,
        "maxOccurs": 1,
        "pathInDataset": "/rootName/.../nodeName2"
      }
    ]
  },
  "content": {
    "nodeName1": {
      "content": true
    },
    "nodeName2": {
      "content": -5,
      "validation": [
        {
          "level": "error",
          "message": "Value must be greater than or equal to 0."
        }
      ]
    }
  ]
}
}
}

```

See also [Select operation](#) [p 671]

- **Table**

JSON object containing the properties:

- (Optional) The [table meta data](#) [p 706],
- (Optional) The sort criteria applied,
- (Optional) The table validation report,
- rows containing a table of selected records. Each record is represented by an object, if no record is selected then the table is empty.
- (Optional) pagination containing [pagination](#) [p 715] information if activated.

```

{
  "rows": [
    {
      "label": "Claude Levi-Strauss",
      "details": "http://.../root/individu/1",
      "content": {
        "id": {
          "content": 1
        }
      },
      ...
    },
    {
      "label": "Sigmoud Freud",
      "details": "http://.../root/individu/5",
      "content": {
        "id": {
          "content": 2
        }
      },
      ...
    },
    ...
  ]
}

```

```

        "content": 30
      },
      ...
    }
  ],
  "sortCriteria": [
    {
      "path": "/name",
      "order": "lasc"
    },
    ...
  ],
  "pagination": {
    "firstPage": null,
    "previousPage": null,
    "nextPage": "http://.../root/individu?pageRecordFilter=./id=9&pageSize=9&pageAction=next",
    "lastPage": "http://.../root/individu?pageSize=9&pageAction=last"
  }
}

```

See also [Select operation](#) [p 671]

- **Record**

JSON object containing:

- The label,
- (Optional) The record URL,
- (Optional) The [technical data](#) [p 718],
- (Optional) The [table metadata](#) [p 706],
- (Optional) The record validation report,
- (Optional) The inheritance mode of the record is: root, inherit, overwrite or occult. This value is available for a child dataset,

See also

[Record lookup mechanism](#) [p 270]

[Inheritance](#) [p 692]

- The record content.

```

{
  "label": "Name1",
  "details": "http://.../rootName/table1/pk1",
  "creationDate": "2015-02-02T19:00:53.142",
  "creationUser": "admin",
  "lastUpdateDate": "2015-09-01T17:22:24.684",
  "lastUpdateUser": "admin",
  "inheritanceMode": "root",
  "meta": {
    "name": "table1",
    "label": "Table1 localized label",
    "type": "table",
    "minOccurs": 0,
    "maxOccurs": "unbounded",
    "primaryKeys": [
      "/pk"
    ],
    "inheritance": "true",
    "fields": [
      {
        "name": "pk",
        "label": "Identifier",
        "type": "string",
        "minOccurs": 1,
        "maxOccurs": 1,
        "pathInRecord": "pk",
        "filterable": true,
        "sortable": true
      }
    ]
  }
}

```



```

    {
      "name": "name",
      "label": "Name",
      "type": "string",
      "minOccurs": 1,
      "maxOccurs": 1,
      "pathInRecord": "name",
      "filterable": true,
      "sortable": true
    },
    {
      "name": "name-fr",
      "label": "Nom",
      "type": "string",
      "minOccurs": 1,
      "maxOccurs": 1,
      "inheritedField": {
        "sourceNode": "../name"
      },
      "pathInRecord": "name-fr",
      "filterable": true,
      "sortable": true
    },
    {
      "name": "parent",
      "label": "Parent",
      "description": "Localized description.",
      "type": "foreignKey",
      "minOccurs": 1,
      "maxOccurs": 1,
      "foreignKey": {
        "tablePath": "/rootName/table1",
        "details": "http://.../rootName/table1"
      },
      "enumeration": "foreignKey",
      "pathInRecord": "parent",
      "filterable": true,
      "sortable": true
    }
  ],
  "content": {
    "pk": {
      "content": "pk1"
    },
    "name": {
      "content": "Name1"
    },
    "name-fr": {
      "content": "Name1",
      "inheritedFieldMode": "inherit"
    },
    "parent": {
      "content": null,
      "validation": [
        {
          "level": "error",
          "message": "Field 'Parent' is mandatory."
        }
      ]
    }
  },
  "validation": {
    ...
  }
}

```

See also [Select operation](#) [p 671]

- **Fields**

For association or selection nodes, contains the target table with associated records if, and only if, the `includedDetails` parameter is set to `true`.

For other kinds of nodes, contains the current [node value](#) [p 711], by adding `meta` entry if enabled.

See also [Select operation](#) [p 671]

- **Retrieve the user interface state**

Contains the user interface status and the unavailability message.

```
{
  "content": {
    "toolStatus": {
      "content": true,
      "label": "Open",
      "selector": "http://.../domain/toolStatus/toolStatus?selector=true",
    },
    "toolStatusCloseMessage": {
      "content": "Access is temporarily forbidden for maintenance.",
    }
  }
}
```

See also [User interface operations](#) [p 667]

**Note**

Node, records and field in meta, rows and content may be hidden depending on their resolved permissions (see [permissions](#) [p 273]).

## 100.3 Meta-data

This section can be activated on demand with the [includeMeta](#) [p 674] parameter. It describes the structure and the JSON typing of the content section.

This section is deactivated by default for selection operations.

See also [Select operation](#) [p 671]

## Structure of table

Table meta-data is represented by a JSON object with the following properties:

| JSON property    | JSON format      | Description                                                                                                                                                                                                           | Required |
|------------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| name             | String           | Name of the table defined in schema.                                                                                                                                                                                  | Yes      |
| label            | String           | Table label. If undefined, the name of the schema node is returned.                                                                                                                                                   | Yes      |
| description      | String           | Table description.                                                                                                                                                                                                    | No       |
| type             | String           | Always equal to: table.                                                                                                                                                                                               | Yes      |
| minOccurs        | Number           | Number of minimum authorized record(s).                                                                                                                                                                               | Yes      |
| maxOccurs        | Number or String | Number of maximum authorized record(s) or unbounded.                                                                                                                                                                  | Yes      |
| history          | Boolean          | Specifies if the table content is historized. Its value is true if history is activated, false otherwise.<br><br><i>See also <a href="#">History</a> [p 249]</i>                                                      | No       |
| primaryKeyFields | Array            | Array of the paths corresponding to the primary key.                                                                                                                                                                  | Yes      |
| inheritance      | Boolean          | Specifies whether the dataset inheritance is activated for the table. Its value is true if inheritance is activated, false otherwise.<br><br><i>See also <a href="#">Inheritance and value resolution</a> [p 268]</i> | No       |
| fields           | Array            | Array of fields, that are direct children of the record node. Each field may also recursively contain sub-fields.                                                                                                     | Yes      |

## Structure of field

Each authorized field is represented by a JSON object with the following properties:

| JSON property  | JSON format      | Description                                                                                                                                                                                                                                                            | Required |
|----------------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| name           | String           | Name of the current authorized schema node.                                                                                                                                                                                                                            | Yes      |
| label          | String           | Node label. If undefined, the name of the schema node is returned.                                                                                                                                                                                                     | Yes      |
| description    | String           | Node description.                                                                                                                                                                                                                                                      | No       |
| type           | String           | Node type: <a href="#">simple type</a> [p 713], group, table, foreignKey, etc.                                                                                                                                                                                         | Yes      |
| minOccurs      | Number           | Number of minimum authorized occurrence(s).                                                                                                                                                                                                                            | Yes      |
| maxOccurs      | Number or String | Number of maximum authorized occurrence(s) or unbounded.                                                                                                                                                                                                               | Yes      |
| inheritedField | Object           | <p>Holds information related to the inherited field's value source.</p> <pre>"inheritedField": {   "sourceRecord": "/path/to/record", // (optional)   "sourceNode": ". /path/to/Node" }</pre> <p>See also <a href="#">Inheritance and value resolution</a> [p 268]</p> | No       |
| foreignKey     | Object           | <p>Contains information related to the target table.</p> <pre>{   "dataspace": "BAuthors",   "dataset": "Authors",   "tablePath": "/root/Authors",   "details": "http://.../BAuthors/Authors/root/Authors" }</pre>                                                     | No (*)   |
| dataspace      | String           | Target dataspace or snapshot identifier.<br>This property is placed under the foreignKey property.                                                                                                                                                                     | No (*)   |
| dataset        | String           | Target dataset identifier.<br>This property is placed under the foreignKey property.                                                                                                                                                                                   | No (*)   |
| tablePath      | String           | Target table path.<br>This property is placed under the foreignKey property.                                                                                                                                                                                           | Yes      |
| details        | String           | Target table URL.<br>This property is placed under the foreignKey property and is included by default.                                                                                                                                                                 | No       |
| enumeration    | String           | Specifies if the field is an enumeration value. Possible values are: <ul style="list-style-type: none"> <li>foreignKey</li> </ul>                                                                                                                                      | No       |

| JSON property | JSON format              | Description                                                                                                                                                                                                                                                                                                                | Required |
|---------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
|               |                          | <ul style="list-style-type: none"> <li>static</li> <li>dynamic</li> <li>programmatic</li> <li>nomenclature</li> <li>resource</li> </ul> <p>See also <i>SchemaFacetEnumeration</i><sup>API</sup></p> <p>Specifies whether the field can be used for retrieving possible values by using the selector request parameter.</p> |          |
| valueFunction | Boolean                  | <p>Specifies if the field is a computed value.</p> <p>See also <a href="#">Computed values</a> [p 529]</p>                                                                                                                                                                                                                 | No       |
| pathInDataset | String                   | Relative field path starting from the schema node.                                                                                                                                                                                                                                                                         | No (**)  |
| pathInRecord  | String                   | Relative field path starting from the table node.                                                                                                                                                                                                                                                                          | No (*)   |
| filterable    | Boolean                  | Specifies whether the field can be used for filtering record using filter request parameter.                                                                                                                                                                                                                               | No (*)   |
| sortable      | Boolean                  | Specifies whether the field can be used in sort criteria using sort request parameter.                                                                                                                                                                                                                                     | No (*)   |
| fields        | Array of Object elements | Contains the structure and typing of each group field.                                                                                                                                                                                                                                                                     | No       |

(\*) Only available for table, record and record field operations.

(\*\*) Only available for dataset tree operations.

## 100.4 Sort criteria information

The sort criteria applied to the request can be returned on demand, by using the [includeSortCriteria](#) [p 675] parameter (deactivated by default). If it is activated, a `sortCriteria` property is directly added to the response root node.

A `sortCriteria` property contains a JSON Array that contains ordered sort criteria, and for each sort criterion, a JSON object is added with the following properties:

| JSON property | JSON format | Description                                    | Required |
|---------------|-------------|------------------------------------------------|----------|
| path          | String      | Field path.                                    | Yes      |
| order         | String      | Possible values are: asc, lasc, desc or ldesc. | Yes      |

## 100.5 Validation

The validation can be activated on demand with the [includeValidation](#) [p 675] parameter (deactivated by default). If it is activated, validation properties are directly added on target nodes with one or several messages. For messages without a target node path, a validation property is added on the root node.

A validation property contains a JSON Array and for each message, corresponding to a validation item, a JSON object with properties:

| JSON property | JSON format                                 | Description                                                                                                                                                                                                                                  | Required |
|---------------|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| level         | String                                      | Severity of the validation item, the possible values are: info, warn, error.                                                                                                                                                                 | Yes      |
| message       | String                                      | Description of the validation item.                                                                                                                                                                                                          | Yes      |
| details       | String<br>corresponding to an absolute URL. | URL of the resource associated with the validation item.<br><br>Only available on the table and dataset scopes, if associated resources exist and if it is included.<br><br><i>See also <a href="#">includeDetails parameter</a> [p 674]</i> | No       |

## 100.6 Content

This section can be deactivated on demand with the [includeContent](#) [p 674] parameter (activated by default). It provides the content of the record values, dataset, or field of one of the content fields for an authorized user. It also has additional information, including labels, technical information, URLs...

The content is represented by a JSON object with a property set for each sub-node.

## Node value

| JSON property      | JSON format                                                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Required |
|--------------------|-----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| content            | <a href="#">Content of simple type</a> [p 713]<br><a href="#">Content of group and list</a> [p 715] | Contains the node value. Available for all nodes except association and selection. However, their content can be retrieved by invoking the URL provided in details.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | No       |
| details            | String<br>corresponding to an absolute URL.                                                         | By invocation, the node details are returned.<br>Response type after invocation depending on the meta type. <ul style="list-style-type: none"> <li>foreignKey: target record (available on table, record and field operation).</li> <li>resource: target <a href="#">resource</a> [p 520] (available on dataset node, table, record and field operations).</li> <li>association: target table containing associated records (available on table and record operations).</li> <li>selection: target table containing associated records (available on table and record operations).</li> <li>group: target dataset group node (available on dataset tree operation).</li> <li>table: target table (available on dataset tree operation).</li> </ul> Example:<br>http://.../BReference/dataset/root/table/pk/associationField | No       |
| label              | String                                                                                              | Contains the foreign key or enumeration label in the current locale.<br>The default label is returned if the current locale is not supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No       |
| inheritanceMode    | String                                                                                              | Contains the node's inheritance state, considering only dataset inheritance. <code>inheritedFieldMode</code> and <code>inheritanceMode</code> properties cannot be both defined on the same node.<br><br>See also<br><a href="#">inheritedFieldMode</a> [p 711]<br><a href="#">Record lookup mechanism</a> [p 270]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | No       |
| inheritedFieldMode | String                                                                                              | Contains the node's field inheritance state, considering dataset and field inheritance. When both inheritances are used, field inheritance has priority over the dataset one. <code>inheritedFieldMode</code> and <code>inheritanceMode</code> properties cannot be both defined on the same node.<br><br>See also<br><a href="#">inheritanceMode</a> [p 711]<br><a href="#">Value lookup mechanism</a> [p 271]                                                                                                                                                                                                                                                                                                                                                                                                             | No       |
| selector           | String                                                                                              | Correspond to the URL for <a href="#">selector</a> [p 716] operation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | No       |

| JSON property | JSON format                       | Description                                                                                       | Required |
|---------------|-----------------------------------|---------------------------------------------------------------------------------------------------|----------|
|               | corresponding to an absolute URL. | Example:<br><code>http://.../BReference/dataset/root/table/pk/<br/>enumField?selector=true</code> |          |
| validation    | Array                             | Contains the validation report that concerns the current node context.                            | No       |



## Content of simple type

A simple field value is stored in a JSON object and the content is the value of the content property.

| XML Schema                                                | JSON format                                                                                                                    | Examples                                                                                                                                                                      | Meta type                               |
|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| xs:string<br>xs:Name<br>osd:html<br>osd:email<br>osd:text | String (Unicode characters, cf. <a href="#">RFC4627</a> )                                                                      | "A text"<br>"The escape of \"special character\" is preceded by a backslash."<br>"<p>An HTML tag can thus be written without trouble</p>"<br>"employee@mycompany.com"<br>null | string<br>name<br>html<br>email<br>text |
| osd:locale                                                | String (Language tag, cf. <a href="#">RFC1766</a> )                                                                            | "en-US"                                                                                                                                                                       | locale                                  |
| xs:string<br>(Foreign key)                                | String<br>contains the value of the formatted foreign key.                                                                     | "0"<br>"true 99"                                                                                                                                                              | foreignKey                              |
| xs:boolean                                                | Boolean                                                                                                                        | true<br>false<br>null                                                                                                                                                         | boolean                                 |
| xs:decimal                                                | Number or null                                                                                                                 | -10.5<br>20.001<br>15<br>-1e-13                                                                                                                                               | decimal                                 |
| xs:date                                                   | String with format: "yyyy-MM-dd"                                                                                               | "2015-04-13"                                                                                                                                                                  | date                                    |
| xs:time                                                   | String with format: <ul style="list-style-type: none"> <li>"HH:mm:ss"</li> <li>"HH:mm:ss.SSS"</li> </ul>                       | "11:55:00"<br>"11:55:00.000"                                                                                                                                                  | time                                    |
| xs:dateTime                                               | String with format: <ul style="list-style-type: none"> <li>"yyyy-MM-ddTHH:mm:ss"</li> <li>"yyyy-MM-ddTHH:mm:ss.SSS"</li> </ul> | "2015-04-13T11:55:00"<br>"2015-04-13T11:55:00.000"                                                                                                                            | dateTime                                |
| xs:anyURI                                                 | String (Uniform Resource Identifier, cf. <a href="#">RFC3986</a> )                                                             | "https://fr.wikipedia.org/wiki/René_Descartes"                                                                                                                                | anyURI                                  |
| xs:int<br>xs:integer                                      | Number or null                                                                                                                 | 1596                                                                                                                                                                          | int                                     |
| osd:resource                                              | String                                                                                                                         | "ebx-tutorial:ext-images:frontpages/Ajax for Dummies.jpg"                                                                                                                     | resource                                |

| XML Schema       | JSON format                                                                                                               | Examples        | Meta type |
|------------------|---------------------------------------------------------------------------------------------------------------------------|-----------------|-----------|
|                  | contains the resource formatted value.                                                                                    |                 |           |
| osd:color        | String with format: "#[A-Fa-f0-9]{6}"<br>contains the formatted value for the color.                                      | "#F6E0E0"       | color     |
| osd:datasetName  | String with format: "[a-zA-Z_-][a-zA-Z0-9_]*" and 64 characters max.<br>contains the formatted value of the dataset name. | "ebx-tutorial"  | dataset   |
| osd:dataspaceKey | String with format: "[BV][a-zA-Z0-9_:\-\\]*" and 33 characters max.<br>contains the formatted key value of the dataspace. | "Bebx-tutorial" | dataspace |

## Content of group and list

| XML Schema              | JSON format                                                                                                                                                  | Examples                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Meta type                                                       |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| Group<br>xs:complexType | object<br>Contains a property per sub-node.                                                                                                                  | Example for a simple-occurrence group.<br><pre>{   "road" : {"content" : "11 rue scribe"},   "zipcode" : {"content" : "75009"},   "country" : {"content" : "France"} }</pre>                                                                                                                                                                                                                                                                                                                                                                         | group                                                           |
| List<br>maxOccurs > 1   | Array<br>Contains an array of all field occurrences represented by a JSON Object.<br><br>Each object is represented as a <a href="#">node value</a> [p 711]. | Example for a multi-occurrence field of the xs:int type.<br><pre>[   {"content": 0},   {"content": 1},   {"content": 2},   {"content": 3} ]</pre><br>Example for a multi-occurrence group.<br><pre>[   {     "content":     {       "road": {"content": "11 rue scribe"},       "zipcode": {"content": "75009"},       "country": {"content": "France"}     }   },   {     "content":     {       "road": {"content": "711 Atlantic Ave"},       "zipcode": {"content": "MA 02111"},       "country": {"content": "United States"}     }   } ]</pre> | <a href="#">Meta of simple type</a> [p 713],<br><br>or<br>group |

## Pagination

This feature allows returning a limited and parameterizable number of data. Pagination can be applied to data of the following types: records, association values, selection node values and selectors. A context named `pagination` is returned only if it has been activated. This context allows browsing data similarly to the UI.

Pagination is activated by default.

See also [Select operation](#) [p 671]

Detailed information related to this context can be found hereafter:

| JSON property | JSON format        | Description                      | Required |
|---------------|--------------------|----------------------------------|----------|
| firstPage     | String or null (*) | URL to access the first page.    | Yes (**) |
| previousPage  | String or null (*) | URL to access the previous page. | Yes (**) |
| nextPage      | String or null (*) | URL to access the next page.     | Yes      |
| lastPage      | String or null (*) | URL to access the last page.     | Yes (**) |

#### Note

(\*) Only defines if data is available in this context and not in the response.

#### Note

(\*\*) Not present on selector.

## Selector

By invoking the URL represented by the property selector on a field that provides an enumeration, this returns a JSON object containing the properties:

- rows containing an Array of JSON object where each one contains two entries, such as the returned content that can be persisted and the corresponding label. The list of possible items is established depending on the current context.
- (Optional) pagination containing [pagination](#) [p 715] information (activated by default).

```
{
  "rows": [
    {
      "content": "F",
      "label": "feminine"
    },
    {
      "content": "M",
      "label": "masculine"
    }
  ],
  "pagination": {
    "nextPage": null
  }
}
```

See also [includeSelector](#) [p 675]

## Insert operation report

When invoking the insert operation with a record table, it can optionally return a report. The report includes a JSON object that contains the following properties:

- rows contains a JSON objectArray, where each element corresponds to the result of a request element.
- code contains an int of the JSON Number type, and allows to know whether the record has been inserted or updated. This property is included if, and only if, the updateOrInsert parameter is set to true.

- `foreignKey` contains a string of the JSON string type, corresponding to the content to be used as a foreign key for this record. This property is included if, and only if, the parameter `includeForeignKey` is set to `true`.
- `label` contains a string of the JSON string type, and allows to retrieve the record label. This property is included if, and only if, the parameter `includeLabel` is set to `yes`.
- `details` containing a string of the JSON string type, corresponding to the resource URL. This property is included if, and only if, the parameter `includeDetails` is set to `true`.

```
{
  "rows": [
    {
      "code": 204,
      "foreignKey": "62",
      "label": "Claude Debussy",
      "details": "http://.../root/individu/62"
    },
    {
      "code": 201,
      "foreignKey": "195",
      "label": "Camille Saint-Saëns",
      "details": "http://.../root/individu/195"
    }
  ]
}
```

See also [Insert operation](#) [p 679]

## Delete operation report

When invoking the delete operation, a report is returned. The report includes a JSON object that contains the following properties:

- `deletedCount` containing an integer of the JSON Number type, corresponds to the number of deleted records.
- `occultedCount` containing an integer of the JSON Number type, corresponds to the number of occulted records.
- `inheritedCount` containing an integer of the JSON Number type, corresponds to the number of inherited records.

```
{
  "deletedCount": 1,
  "occultedCount": 0,
  "inheritedCount": 0
}
```

See also [Delete operation](#) [p 686]

## Technical data

Each returned record is completed with the properties corresponding to its technical data, containing:

| JSON property  | JSON format | Description                  | Required |
|----------------|-------------|------------------------------|----------|
| creationDate   | String      | Creation date.               | Yes      |
| creationUser   | String      | Creation user identifier.    | Yes      |
| lastUpdateDate | String      | Last update date.            | Yes      |
| lastUpdateUser | String      | Last update user identifier. | Yes      |

```
{
  ...
  "creationDate": "2015-12-24T19:00:53.158",
  "creationUser": "admin",
  "lastUpdateDate": "2015-12-25T00:00:00.001",
  "lastUpdateUser": "admin",
  ...
}
```

## 100.7 Update modes

The byDelta mode allows to ignore data model elements that are missing from the JSON source document. This mode is enabled (by default) through RESTful operations. The following table summarizes the behavior of insert and update operations when elements are not included in the source document.

See the RESTful data services operations [update](#) [p 684] and [insert](#) [p 679], as well as `ImportSpec.setByDeltaAPI` in the Java API for more information.

| State in the JSON source document                                                                      | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The property does not exist in the source document                                                     | <p><b>If the byDelta mode is activated (default):</b></p> <ul style="list-style-type: none"> <li>For the update operation, the field value remains unchanged.</li> <li>For the insert operation, the behavior is the same as when the byDelta mode is disabled.</li> </ul> <p><b>If the byDelta mode is disabled through the RESTful operation parameter:</b></p> <p>The target field is set to one of the following values:</p> <ul style="list-style-type: none"> <li>If the element defines a default value, the target field is set to that default value.</li> <li>If the element is of a type other than a string or list, the target field value is set to <code>null</code>.</li> <li>If the element is an aggregated list, the target field value is set to an empty list value.</li> <li>If the element is a string that differentiates <code>null</code> from an empty string, the target field value is set to <code>null</code>. If it is a string that does not differentiate the two, an empty string.</li> <li>If the element (simple or complex) is hidden in the data services, the target value remains unchanged.</li> </ul> <p>See also <a href="#">Hiding a field in Data Services</a> [p 542]</p> <p><b>Note</b></p> <p>The user performing the import must have the required permissions to create or change the target field value. Otherwise, the operation will be aborted.</p> |
| The element is present and its value is <code>null</code> (for example, <code>"content": null</code> ) | The target field is always set to <code>null</code> except for lists, in which case it is not supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## 100.8 Known limitations

### **Field values**

The value of fields `xs:date`, `xs:time` and `xs:dateTime` does not contain a time zone associated with the JSON-primitive type.





# CHAPTER 101

---

## REST Toolkit

This chapter contains the following topics:

1. [Introduction](#)
2. [Application definitions](#)
3. [Service and operation definitions](#)
4. [Authentication and lookup mechanism](#)
5. [REST authentication and permissions](#)
6. [URI builders](#)
7. [Exception handling](#)
8. [Monitoring](#)
9. [Packaging and registration](#)

### 101.1 Introduction

TIBCO EBX® offers the possibility to develop custom REST services using the REST Toolkit. The REST Toolkit supports JAX-RS 2.1 ([JSR-370](#)) and JSON-B ([JSR-367](#)).

A REST service is implemented by a Java class and its operations are implemented by Java methods. The response can be generated by serializing POJO objects. The request input can be unserialized to POJOs. Various input and output formats, including JSON, are supported. For more details on supported formats, see [media types](#) [p 723].

Rest Toolkit supports the following:

- **Injectable objects**  
EBX® provides injectable objects useful to authenticate the request's user, to access the EBX® repository or to build URIs without worrying about the configuration (for example [reverse-proxy](#) [p 652] or [REST forward](#) [p 580] modes);  
JAX-RS injectable objects are also supported.
- **Annotations**  
EBX® provides annotations to describe resources, grant anonymous access or add restrictions to a method.  
JAX-RS and JSON-B annotations are also supported.
- **logging and debugging utilities.**

See also [JAX-RS: Java™ API for RESTful Web Services 2.1](#)

## 101.2 Application definitions

An EBX® module, that includes custom REST services, must provide at least one REST Toolkit application class. A REST Toolkit application class extends the EBX® `RESTApplicationAbstract`<sup>API</sup> class. The minimum requirement is to define the base URL, using the `@ApplicationPath` annotation and the set of packages to scan for REST service classes.

### Note

Only packages accessible from the web application's classloader can be scanned.

### Note

It is possible to register REST resource classes or singletons, packaged inside or outside the web application archive, through the **`ApplicationConfigurator.register(java.lang.Class)`** `ApplicationConfigurator.register`<sup>API</sup> or **`ApplicationConfigurator.register(java.lang.Object)`** `ApplicationConfigurator.register`<sup>API</sup> methods.

### Note

If no packages scope is defined, then every class reachable from the web application's classloader will be scanned.

The application path cannot be "/" and must not collide with an existing resource from the module. It is recommended to use "/rest" (the value of the `RESTApplicationAbstract.REST_DEFAULT_APPLICATION_PATH` constant).

EBX® `Documentation`<sup>API</sup> annotation is optional. It is displayed to administrators in 'Technical configuration' > 'Modules and data models' or when logging and debugging.

```
import javax.ws.rs.*;

import com.orchestranetworks.rest.*;
import com.orchestranetworks.rest.annotation.*;

@ApplicationPath(RESTApplicationAbstract.REST_DEFAULT_APPLICATION_PATH)
@Documentation("My REST sample application")
public final class RESTApplication extends RESTApplicationAbstract
{
    public RESTApplication()
    {
        // Adds one or more package names which will be used to scan for components.
        super((cfg) -> cfg.addPackages(RESTApplication.class.getPackage()));
    }
}
```

## 101.3 Service and operation definitions

A REST Toolkit service is implemented by a Java class and its operations are implemented by its methods.

Class and methods can be annotated by `@Path` to specify the access path. The `@Path` annotation value defined at the class level will prepend the ones defined on methods. Warning, only one `@Path` annotation is allowed per class or method.

Media types accepted and produced by a resource are respectively defined by the `@Consumes` and `@Produces` JAX-RS annotations. The supported media types are:

- `application/json` ( [MediaType.APPLICATION\\_JSON\\_TYPE](#) )
- `application/octet-stream` ( [MediaType.APPLICATION\\_OCTET\\_STREAM\\_TYPE](#) )
- `application/x-www-form-urlencoded` ( [MediaType.APPLICATION\\_FORM\\_URLENCODED\\_TYPE](#) )
- `multipart/form-data` ( [MediaType.MULTIPART\\_FORM\\_DATA\\_TYPE](#) )
- `text/css`
- `text/html` ( [MediaType.TEXT\\_HTML\\_TYPE](#) )
- `text/plain` ( [MediaType.TEXT\\_PLAIN\\_TYPE](#) )

Valid HTTP(S) methods are specified by JAX-RS annotations `@GET`, `@POST`, `@PUT`, etc. Only one of these annotations can be set on each Java method (this means that a Java method can support only one HTTP method).

**Warning:** URL parameters with a name prefixed with `ebx-` are reserved by REST Toolkit and should not be defined by custom REST services, unless explicitly authorized by the EBX® documentation.

## URL and sample

The REST URL to access the description service for the sample is defined below:

`http[s]://<host>[:<port>]/<path to webapp>/rest/track/v1/description`

Where:

- `<path to webapp>` corresponds to the web application's path holding the REST Toolkit application, itself serving the sample service. The path is composed by multiple, or none, URI segments followed by the web application's name.

### Note

Please note that `/rest/track/v1/description` corresponds to the concatenation of the application's `@ApplicationPath` and service's `@Path` annotations.

The following REST Toolkit service sample provides methods to query and manage track data:

```
import java.net.*;
import java.util.*;
import java.util.concurrent.*;
import java.util.regex.*;
import java.util.stream.*;

import javax.servlet.http.*;
import javax.ws.rs.*;
import javax.ws.rs.container.*;
import javax.ws.rs.core.*;

import com.orchestranetworks.rest.annotation.*;
import com.orchestranetworks.rest.inject.*;

/**
 * The REST Toolkit Track service v1.
 */
@Produces(MediaType.APPLICATION_JSON)
@Consumes(MediaType.APPLICATION_JSON)
@Path("/track/v1")
@Documentation("Track service")
public final class TrackService
{
    @Context
    private ResourceInfo resourceInfo;
```

```

@Context
private SessionContext sessionContext;

private static final Map<Integer, TrackDTO> TRACKS = new ConcurrentHashMap<>();

/**
 * Gets service description
 */
@GET
@Path("/description")
@Documentation("Gets service description")
@Produces({ MediaType.TEXT_PLAIN, MediaType.APPLICATION_JSON })
@AnonymousAccessEnabled
public String handleServiceDescription()
{
    return this.resourceInfo.getResourceMethod().getAnnotation(Documentation.class).value();
}

/**
 * Selects tracks.
 */
@GET
@Path("/tracks")
@Documentation("Selects tracks")
public Collection<TrackDTO> handleSelectTracks(
    @QueryParam("singerFilter") final String singerFilter, // a URL parameter holding a Java regular expression
    @QueryParam("titleFilter") final String titleFilter) // a URL parameter holding a Java regular expression
{
    final Pattern singerPattern = TrackService.compilePattern(singerFilter);
    final Pattern titlePattern = TrackService.compilePattern(titleFilter);

    return TRACKS.values()
        .parallelStream()
        .filter(Objects::nonNull)
        .filter(track -> singerPattern == null || singerPattern.matcher(track.singer).matches())
        .filter(track -> titlePattern == null || titlePattern.matcher(track.title).matches())
        .collect(Collectors.toList());
}

private static Pattern compilePattern(final String aPattern)
{
    if (aPattern == null || aPattern.isEmpty())
        return null;

    try
    {
        return Pattern.compile(aPattern);
    }
    catch (final PatternSyntaxException ignore)
    {
        // ignore invalid pattern
        return null;
    }
}

/**
 * Counts all tracks.
 */
@GET
@Path("/tracks:count")
@Documentation("Counts all tracks")
public int handleCountTracks()
{
    return TRACKS.size();
}

/**
 * Selects a track by id.
 */
@GET
@Path("/tracks/{id}")
@Documentation("Selects a track by id")
public TrackDTO handleSelectTrackById(@PathParam("id") Integer id)
{
    final TrackDTO track = TRACKS.get(id);
    if (track == null)
        throw new NotFoundException("Track id [" + id + "] does not found.");
    return track;
}

/**
 * Deletes a track by id.
 */
@DELETE
@Path("/tracks/{id}")

```

```

@Documentation("Deletes a track by id")
public void handleDeleteTrackById(@PathParam("id") Integer id)
{
    if (!TRACKS.containsKey(id))
        throw new NotFoundException("Track id [" + id + "] does not found.");
    TRACKS.remove(id);
}

/**
 * Inserts or updates one or several tracks.
 * <p>
 * The complex response structure corresponds to one of:
 * <ul>
 * <li>An empty content with the <code>location</code> HTTP header defined
 *     to the access URI.</li>
 * <li>A JSON array of {@link ResultDetailsDTO} objects.</li>
 * </ul>
 */
@POST
@Path("/tracks")
@Documentation("Inserts or updates one or several tracks")
public Response handleInsertOrUpdateTracks(List<TrackDTO> tracks)
{
    int inserted = 0;
    int updated = 0;

    final ResultDetailsDTO[] resultDetails = new ResultDetailsDTO[tracks.size()];
    int resultIndex = 0;

    final URI base = this.sessionContext.getUriInfoUtility()
        .createBuilderForRESTApplication()
        .path(this.getClass())
        .segment("tracks")
        .build();

    for (final TrackDTO track : tracks)
    {
        final String id = String.valueOf(track.id);
        final URI uri = UriBuilder.fromUri(base).segment(id).build();

        final int code;
        if (TRACKS.containsKey(track.id))
        {
            code = HttpServletResponse.SC_NO_CONTENT;
            updated++;
        }
        else
        {
            code = HttpServletResponse.SC_CREATED;
            inserted++;
        }

        TRACKS.put(track.id, track);

        resultDetails[resultIndex++] = ResultDetailsDTO.create(
            code,
            null,
            String.valueOf(track.id),
            uri);
    }

    if (inserted == 1 && updated == 0)
        return Response.created(resultDetails[0].details).build();

    return Response.ok().entity(resultDetails).build();
}

/**
 * Updates one track.
 */
@PUT
@Path("/tracks/{id}")
@Documentation("Update one track")
public void handleUpdateOneTrack(@PathParam("id") Integer id, TrackDTO aTrack)
{
    final TrackDTO track = TRACKS.get(id);
    if (track == null)
        throw new NotFoundException("Track id [" + id + "] does not found.");

    if (aTrack.id != null && !aTrack.id.equals(track.id))
        throw new BadRequestException("Selected track id [" + id
            + "] is not equals to body track id.");

    TRACKS.put(aTrack.id, aTrack);
}

```

```
}
```

This REST service uses the following Java classes, which represent a Data Transfer Objects (DTO), to serialize and deserialize data:

```
/**
 * DTO for a track.
 */
public final class TrackDTO
{
    public Integer id;
    public String singer;
    public String title;
}

import java.net.*;

/**
 * DTO for result details.
 */
@JsonbPropertyOrder({ "code", "label", "foreignKey", "details" })
public final class ResultDetailsDTO
{
    public int code;
    public String label;
    public String foreignKey;
    public URI details;

    public static ResultDetailsDTO create(
        final int aCode,
        final String aForeignKey,
        final URI aDetails)
    {
        return ResultDetailsDTO.create(aCode, null, aForeignKey, aDetails);
    }

    public static ResultDetailsDTO create(
        final int aCode,
        final String aLabel,
        final String aForeignKey,
        final URI aDetails)
    {
        final ResultDetailsDTO result = new ResultDetailsDTO();
        result.code = aCode;
        result.label = aLabel;
        result.foreignKey = aForeignKey;
        result.details = aDetails;
        return result;
    }
}
```

## 101.4 Authentication and lookup mechanism

A custom REST service developed with REST Toolkit supports the same authentication methods and lookup mechanism as the built-in REST data services. However, there is a slight difference concerning the 'Anonymous authentication Scheme' since its scope can be wider by using the `AnonymousAccessEnabledAPT`. See [REST authentication and permissions](#) [p 726] for more information.

See also

[Authentication](#) [p 654]

[Lookup mechanism](#) [p 655]

## 101.5 REST authentication and permissions

By default, every REST resource Java method requires users to be authenticated.

However, some methods may need to be accessible anonymously. These methods must be annotated by `AnonymousAccessEnabledAPT`.

Some methods may need to be restricted to given profiles. These methods may be annotated by `AuthorizationAPI` to specify an authorization rule. An authorization rule is a Java class that implements the `AuthorizationRuleAPI` interface.

```
import javax.ws.rs.*;

import com.orchestranetworks.rest.annotation.*;

/**
 * The REST Toolkit service v1.
 */
@Path("/service/v1")
@Documentation("Service")
public final class Service
{
    ...

    /**
     * Gets service description
     */
    @GET
    @AnonymousAccessEnabled
    public String handleServiceDescription()
    {
        ...
    }

    /**
     * Gets restricted service
     */
    @GET
    @Authorization(IsUserAuthorized.class)
    public RestrictedServiceDTO handleRestrictedService()
    {
        ...
    }
}
```

## 101.6 URI builders

REST Toolkit provides an utility interface `URIInfoUtilityAPI` to generate URIs. An instance of this interface is accessible through the injectable built-in object `SessionContextAPI`.

## 101.7 Exception handling

A REST Toolkit Java method can produce a standard HTTP error response by throwing a Java exception that extends the JAX-RS class `javax.ws.rs.WebApplicationException`. JAX-RS defines exceptions for various HTTP status codes. EBX® defines `UnprocessableEntityExceptionAPI` that adds support for the HTTP 422(*Unprocessable entity*) code.

Plain Java exceptions are mapped to the HTTP status code 500 (*Internal server error*).

```
{
  "userMessage": "...", // Mandatory localized message
  "errorCode": "999", // EBX® error code (optional, used mainly for HTTP error 422)
  "errors": [
    // Internal messages useful when debugging (optional).
    // Usually not displayed to the user and not localized.
    "Message 1", "Message 2" ]
}
```

## 101.8 Monitoring

REST Toolkit events monitoring is similar to the data services log configuration. The difference is the property key which must be `ebx.log4j.category.log.restServices`.

**See also**[Monitoring](#) [p 655][Configuring the EBX® logs](#) [p 354]

Some additional properties are available to configure the log messages. See [Configuring REST toolkit services](#) [p 359] for further information.

## 101.9 Packaging and registration

All applications and components are required to be packaged into the module's Web Application (war file).

The JAX-RS libraries, except the JAX-RS client API, are already included in `ebx.jar` and must not be packaged in the war file.

See [Java EE deployment](#) [p 315] for more information.

The registration of a REST Toolkit application is integrated into the EBX® module registration process. The registration class must extend `ModuleRegistrationListenerAPI`, declare the Servlet 3.0 annotation `@WebListener` and override the `handleContextInitialized` method.

See [Module registration](#) [p 462] for more information.

```
import javax.servlet.annotation.*;

import com.orchestranetworks.module.*;

@WebListener
public final class RegistrationModule extends ModuleRegistrationListener
{
    @Override
    public void handleContextInitialized(final ModuleInitializedContext aContext)
    {
        // Registers dynamically a REST Toolkit application.
        aContext.registerRESTApplication(RESTApplication.class);
    }
}
```