

User Guide

- Introduction
 - [How TIBCO EBX® works](#)
 - [Using the user interface](#)
 - [Glossary](#)
- Data models
 - [Introduction to data models](#)
 - [Implementing data models](#)
 - [Creating a data model](#)
 - [Configuring the data model](#)
 - [Implementing the data model structure](#)
 - [Properties of data model elements](#)
 - [Data validation controls on elements](#)
 - [Data model extensions](#)
 - [Working with an existing data model](#)
 - [Publishing and versioning data models](#)
 - [Publishing a data model](#)
 - [Versioning an embedded data model](#)
- Dataspaces
 - [Introduction to dataspaces](#)
 - [Creating a dataspace](#)
 - [Working with existing dataspaces](#)
 - [Snapshots](#)
- Datasets
 - [Introduction to datasets](#)

- Creating a dataset
- Viewing data
- Editing data
- File import and export services
 - XML import and export
 - CSV import and export
- Working with existing datasets
- Dataset inheritance
- Collaborative features (Team Up)
 - Introduction to collaborative features (Team Up)
 - Comments
 - Ratings
 - Tags
- Workflow models
 - Introduction to workflow models
 - Creating and implementing a workflow model
 - Configuring the workflow model
 - Publishing workflow models
- Data workflows
 - Introduction to data workflows
 - Using the Data Workflows area user interface
 - Work items
 - Managing data workflows
 - Launching and monitoring data workflows
 - Administration of data workflows
- Data services
 - Introduction to data services

- [Generating data service WSDLs](#)

Introduction

- [How TIBCO EBX® works](#)
- [Using the user interface](#)
- [Glossary](#)

How TIBCO EBX® works

- [Product overview](#)
- [EBX® architecture](#)

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](#)
- [Datasets](#)
- [Dataspaces](#)
- [Workflow modeling](#)
- [Data workflows](#)
- [Data services](#)

EBX® architecture

The following diagram illustrates the EBX® architecture.



Using the user interface

- [Overview](#)
- [Advanced perspective](#)
- [Perspectives](#)
 - [Perspective modes](#)
- [User pane](#)
 - [Avatar](#)
 - [Density](#)
- [User interface features](#)
 - [Resetting the navigation pane width](#)
- [Where to find EBX® help](#)

- [Context-sensitive help](#)
- [Contextual help on elements](#)

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](#)

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*.




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](#).

Example of a hierarchical menu:



Perspective modes


It's possible now to choose between two modes of perspective.

According to the mode chosen by the user, the last visited perspective or the favorite one will be applied at connection.

Last used perspective:

This option allows the user to get the last selected perspective and it will be applied in the next log in.


If the user chooses this mode, the stars next to the perspective name will disappear.



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](#)

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.addUserAvatar
```

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

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.



Where to find EBX® help


In addition to the full standalone product documentation accessible via the [user pane](#), 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.



Glossary

- Governance
 - repository
 - profile
 - user
 - role
 - administrator
 - user and roles directory
 - user session
- Data modeling
 - data model
 - field
 - primary key
 - foreign key

- table (in data model)
- group
- reusable type
- validation rule
- data model assistant (DMA)
- Datasets
 - record
 - table (in dataset)
 - dataset
 - inheritance
 - views
 - recommended view
 - favorite view
- Data management life cycle
 - dataspace
 - reference dataspace
 - dataspace merge
 - snapshot
- History
 - historization
 - history profile
 - table history view
 - transaction history view
- Workflow modeling
 - workflow model
 - script task
 - user task
 - workflow condition

- sub-workflow invocation
- wait task
- data context
- Data workflows
 - workflow publication
 - data workflow
 - work list
 - work item
 - token
- Data services
 - data service
 - lineage
- Cross-domain
 - node

Governance

repository

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

See also [dataspace](#).

profile

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

See also [user](#), [role](#).

Related Java API `Profile`.

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](#), [profile](#).

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

Related Java API `UserReference`.

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](#), [profile](#).

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

Related Java API `Role`.

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](#), [role](#).

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

Related Java API `Directory`, `DirectoryHandler`.

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](#).

Related Java API `Session`.

Data modeling

Main documentation section [Data models](#)

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](#).

Related concept [Data models](#).

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](#), [group](#), [table \(in data model\)](#), [validation rule](#), [inheritance](#).

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

Related Java API `SchemaNode`.

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](#).

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](#).

Related concept [Foreign key](#).

table (in data model)

A data model element that is composed 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](#), [primary key](#), [reusable type](#).

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](#).

Related Java API `SchemaNode`.

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](#).

Datasets

Main documentation section [Datasets](#)

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 \(in dataset\)](#), [primary key](#).

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](#), [primary key](#).

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 \(in dataset\)](#), [field](#), [group](#), [views](#).

Related concept [Datasets](#).

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](#).

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, hierarchical or tile 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](#)
- [Hierarchies](#)

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](#).

favorite view

When displaying a table, the user can choose to define the current as their favorite view through the 'View' menu toolbar.

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

Related concept ['View' menu toolbar](#).

Data management life cycle

Main documentation section [Dataspaces](#)

dataspace

A container entity composed 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](#), [repository](#), [dataspace merge](#).

Related concept [Dataspaces](#).

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](#), [dataspace merge](#), [repository](#).

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](#).

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](#)

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

History

Main documentation section [History](#)

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](#), [transaction history view](#), [history profile](#).

history profile

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

See also [history profile](#).

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](#).

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](#).

Workflow modeling

Main documentation section [Workflow models](#)

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](#).

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](#).

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](#).

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.

Data workflows

Main documentation section [Data workflows](#)

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](#).

Related concept [Data workflows](#).

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](#).

token

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

Data services

Main documentation section [Data services](#)

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](#).

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

Data models

- [Introduction to data models](#)
- [Implementing data models](#)
 - [Creating a data model](#)
 - [Configuring the data model](#)
 - [Implementing the data model structure](#)
 - [Properties of data model elements](#)
 - [Data validation controls on elements](#)
 - [Data model extensions](#)
 - [Working with an existing data model](#)
- [Publishing and versioning data models](#)
 - [Publishing a data model](#)
 - [Versioning an embedded data model](#)

Introduction to data models

- [Overview](#)
 - [What is a data model?](#)
 - [Basic concepts used in data modeling](#)
- [Using the Data Models area user interface](#)
 - [Navigating within the Data Model Assistant](#)
 - [Data model element icons](#)

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


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:

Configuration	The technical configuration of the data model.
Global properties	Defines the global properties of the data model.
Included data models	Defines the data models included in the current model. All types defined in included data models can be reused in the current model.
Component library	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.
Add-ons	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.
Data structure	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.
Simple data types	Simple reusable types defined in the current data model.
Complex data types	Complex reusable types defined in the current data model.
Included simple data types	Simple reusable types defined in an included external data model.
Included complex data types	Complex reusable types defined in an included external data model.
Extensions	Extensions available in the current data model.
Toolbars	The toolbars available to use in the data model.
User services	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)..
Entity mappings	Specifies entity names for the tables in the data model. These names are used as aliases in SQL requests or as suffixes in WSDL operations. This allows to refer to a

	table in SQL requests and data service operations using a unique name instead of its path.
Replications	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.
Ajax components	Defines the available Ajax components in the model.
Java bindings	The bindings specify what Java types have to be generated from the model.

See also:

- [Implementing the data model structure](#)
- [Configuring the data model](#)
- [Reusable types](#)
- [Data model extensions](#)

Data model element icons



field



primary key



foreign key



table



group

Related concepts:

- [Dataspaces](#)
- [Datasets](#)

Implementing data models

- [Creating a data model](#)
- [Configuring the data model](#)
- [Implementing the data model structure](#)
- [Properties of data model elements](#)
- [Data validation controls on elements](#)
- [Data model extensions](#)
- [Working with an existing data model](#)

Creating a data model

Creating a new data model

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

Configuring the data model

- [Information associated with a data model](#)
- [Permissions](#)
- [Data model properties](#)
- [Included data models](#)
- [Add-ons used by the data model](#)

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'** menu for your data model in the navigation pane.

Note:

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

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

Permissions

To define the user permissions on your data model, select 'Permissions' from the **data model** **'Actions'** 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](#).

Data model properties

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

Module name	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.
Module path	Physical location of the module on the server's file system.

Sources location	The source path used when configuring Java components in the 'Component library'. If this path is relative, it will be resolved using the 'Module path' as the base path.
Publication mode	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. See Publication modes for more information. Model path in module: Defines the target file for the data model generation. It must start with '/'.
Dataset inheritance	Specifies whether dataset inheritance is enabled for this data model. Dataset inheritance is disabled by default. See Dataset inheritance for more information.
Documentation	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. See Dynamic labels and descriptions for more information.
Special extensions	Access permissions defined by programmatic rules in a Java class.

Disable auto-increment checks	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. See Auto-incremented values for more information.
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 <code>UserServiceDeclaration</code> for more information.

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](#)

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:

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

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.

- [Common actions and properties](#)
 - [Adding elements to the data model](#)
 - [Deleting elements of the data model](#)
 - [Duplicating existing elements](#)
 - [Moving elements](#)
- [Reusable types](#)
 - [Defining a reusable type](#)

- Using a reusable type
- Including data types defined in other data models
- Data model element creation details
 - Creating fields
 - Creating tables
 - Creating groups
 - Creating primary key fields
 - Creating or defining foreign key fields
 - Creating associations
- Modifying existing elements
 - Removing a field from the primary key

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.

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](#) for more information.

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](#) 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](#) 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.

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](#) in the glossary.

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](#)
- [Basic element properties](#)
 - [Common basic properties](#)
 - [Basic properties for fields](#)
- [Advanced element properties](#)
 - [Common advanced properties](#)
 - [Advanced properties for fields](#)
 - [Advanced properties for tables](#)
 - [Advanced properties for groups](#)

Basic element properties

Common basic properties

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

Information	Additional non-internationalized information associated with the element.
Minimum number of values	Minimum number of values for an element. 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.
Maximum number of values	Maximum number of values for an element. When set to a value greater than '1', the element becomes multi-valued. As primary keys cannot be multi-valued, they must have this property set to '1' or 'undefined'. 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.
Validation rules	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. See Criteria editor for more information.

This can be useful if the validation of the value depends on complex criteria or on the value of other fields.

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 [Constraint on 'null' values](#) for more information.

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.

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 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](#) for more information.

Basic properties for fields

The following basic properties are specific to fields:

Default value	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. See Default value for more information.
Conversion error message	Internationalized messages to display to users when they enter a value that is invalid for the data type of this field.
Computation rule	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. See criteria editor This can be useful if the value depends on other values in the same record, but does not require a programmatic computation. The following limitations exist for computation rules: • Computation rules can only be defined on simple fields inside a table. • Computation rules cannot be defined on fields of type <code>OResource</code> or <code>Password</code>. • Computation rules cannot be defined on selection nodes and primary key fields.

- Computation rules cannot be defined when accessing an element from the validation report.

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.

- Structured search tool

Specifies whether or not the current element is shown in a dataset structured search tool. Setting the property to 'Hidden in structured search' will hide the element in the structured search tool of a dataset. The element will remain searchable in the quick search tool. 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 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 <code>UIWidgetFactory</code> 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 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 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 <code>UIBeanEditor</code> 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 <code>NodeDataTransformer</code> 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 <code>RW</code> in the data model XSD. • 'Read only' corresponds to the mode <code>R-</code> in the data model XSD. • 'Not a dataset node' corresponds to the mode <code>CC</code> in the data model XSD. • 'Non-terminal node' corresponds to the mode <code>--</code> in the data model XSD. See Access properties 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 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 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 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](#) 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](#) in the Developer Guide.

UI bean

See [Common advanced properties](#).

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 `ValueFunction` interface.

Disable validation

Specifies if the constraints defined on the field must be disabled. This property can only be defined on function 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](#).

Access properties

See [Common advanced properties](#).

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:

Start value	Value with which to begin the auto-increment. If this attribute is not specified, the default value is '1'.
Increment step	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'.
Disable auto-increment checks	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 occulating 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 dataspaces 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](#) in the Developer Guide.

Default view

See [Common advanced properties](#).

Node category

See [Common advanced properties](#).

Inherited field

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

Source record	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.
Source element	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](#) in the glossary.

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

Advanced properties for tables

The following advanced properties are specific to tables.

Table

Primary key	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. Each primary key field is denoted by its absolute XPath notation that starts under the table's root element. If there are several elements in the primary key, the list is white-space delimited. For example, <code>"/name /startDate"</code>.
Presentation	Specifies how records are displayed in the user interface of this table in a dataset.
<i>Presentation</i> > Record labeling	Defines the fields to provide the default and localized labels for records in the table. 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</code> interface or the <code>UILabelRendererForHierarchy</code> interface. Attention: 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.
<i>Presentation</i> > Default rendering for groups in forms	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. See Record form: rendering mode for nodes in the Administration Guide. Enabled rendering for groups 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. Default rendering for groups 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 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.
<i>Presentation</i> > Specific rendering of forms	Defines a specific rendering for customizing the record form in a dataset. See <code>UIForm</code> and <code>UserServiceRecordFormFactory</code> for more information.
Toolbars	Defines the toolbars to use in this table. Toolbars can be edited in the <i>Configuration > Toolbars</i> section. Tabular view top: Defines the toolbar to use on top of the default table view. Tabular view row: Defines the toolbar to use on each row of the default table view. Tile view item: Defines the toolbar to use on each item of the tile view. Record top: Defines the toolbar to use in the record form. Hierarchy top: Defines the toolbar to use in the default hierarchy view of the table. See Toolbars for more information.
History	Specifies when historization is to be performed, and the level of guarantee requested. The available history profiles can be edited in Administration > History and logs.

	See History configuration in the repository for more information.
Specific filters	Defines record display filters on the table.
Actions	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.

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](#) in the Developer Guide.

Access properties

See [Common advanced properties](#).

Default view

See [Common advanced properties](#).

Node category

See [Common advanced properties](#).

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](#).

Transformation on export

See [Common advanced properties](#).

Access properties

See [Common advanced properties](#).

Default view

Visibility	See Common advanced properties .
Rendering in forms	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. <i>Tab position</i> 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.

Node category

See [Common advanced properties](#).

Related concepts:

- [Data validation controls on elements](#)

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](#)
- [Simple content validation](#)
- [Advanced content validation](#)
 - [Validation properties](#)

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.

Fixed length	The exact number of characters required for this field.
Minimum length	The minimum number of characters allowed for this field.
Maximum length	The maximum number of characters allowed for this field.
Pattern	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.
Decimal places	The maximum number of decimal places allowed for this field.
Maximum number of digits	The maximum total number of digits allowed for this integer or decimal field.

Enumeration	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.
Greater than [constant]	Defines the minimum value allowed for this field.
Less than [constant]	Defines the maximum value allowed for this field.

See [XML schema supported facets](#).

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](#)

Foreign key constraint	
Table	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.

Mode	Location of the table referenced by the foreign key. 'Default': current data model. 'Other dataset': different dataset, in the same dataspace. 'Other dataspace': dataset in a different dataspace.
Referenced table	XPath expression describing the location of the table. For example, <code>/root/MyTable</code>.
Referenced dataset	Required if the table is located in another dataset. The unique name of the dataset containing the referenced table.
Referenced dataspace	Required if the table is located in another dataspace. The unique name of the dataspace containing the referenced table.
Label	Defines fields to provide the default and localized labels for records in the table. Allows as well to customize the display of the specified label in breadcrumb. 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</code> interface of the Java API. Attention: 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.
Filter	Defines a foreign key filter using an XPath expression. Can also specify a Java class that implements the <code>TableRefFilter</code> interface of the Java API.
Greater than [dynamic]	Defines a field to provide the minimum value allowed for this field.
Less than [dynamic]	Defines a field to provide the maximum value allowed for this field.
Fixed length [dynamic]	Defines a field to provide the exact number of characters required for this field.
Minimum length [dynamic]	Defines a field to provide the minimum number of characters allowed for this field.
Maximum length [dynamic]	Defines a field to provide the maximum number of characters allowed for this field.
Excluded values	Defines a list of values that are not allowed for this field.
Excluded segment	Defines an inclusive range of values that are not allowed for this field. Minimum excluded value: Lowest value not allowed for this field. Maximum excluded value: Highest value not allowed for this field.

Specific constraint (component)	Specifies one or more Java classes that implement the <code>Constraint</code> interface of the Java API. See Programmatic constraints for more information.
Specific enumeration (component)	Specifies a Java class to define an enumeration. The class must define an ordered list of values by implementing the <code>ConstraintEnumeration</code> interface of the Java API.
Enumeration filled by another node	Defines the possible values of this enumeration using a reference to another list or enumeration element.
Dataspace set configuration	Define the dataspace sets 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. Includes Specifies the dataspace sets that can be referenced by this field. Pattern: Specifies a pattern that filters dataspace sets. The pattern is checked against the name of the dataspace sets. Type: Specifies the type of dataspace sets 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.

- 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` 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` 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:

Validation	Defines a localized validation message with a user-defined severity level.
Severity	Defines the severity level of the validation message. Possible values are 'Error', 'Warning', and 'Information'.
Error management policy	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.
Message	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](#)

Data model extensions

- [Extensions used by the data model](#)
- [Indexing and search strategies](#)
- [Business objects](#)

- Defining the structure of a business object model
- Toolbars
 - Defining the structure of a toolbar
 - Deleting elements
 - Duplicating existing elements
 - Moving elements
 - Associate with existing tables
 - Exporting the toolbars
 - Importing toolbars
- Entity mappings
- Replication of data to relational tables
- Node confidentiality

Extensions used by the data model

On any data model, it is possible to specify some *extensions* to be used. These extensions will have the capacity to define additional features around the data model and to enrich the current data model after the publication by adding properties and constraints to the data model elements. Data model extensions are displayed in the navigation pane under the section 'Extensions'.

The following extensions are automatically enabled on new data models:

- Toolbars
- Entity mappings
- Replications
- Indexing and search strategies
- Business objects
- Custom forms
- Functions
- Record permissions
- Node confidentiality

Some extensions are optional and defined in dedicated datasets. These extensions can be disabled or enabled through the user interface.

Optional data model extension can be disabled by selecting the action 'Disable extension' from the menu



to the left of extension to disable.

Note:

When disabling an extension its dedicated dataset is also deleted and it is not possible to recover it after its deletion

From the down arrow



to the left of 'Extensions' entry in the navigation pane, you can enable an extension by selecting the action 'Enable an extension'. This option is only displayed if some extensions are available. When selecting this action a form displays the available data model extensions.

Once a data model extension is activated then it is displayed under the entry 'Extensions' in the navigation pane and can be edited as a regular dataset. That is, the configuration of an extension is embedded in a dedicated dataset that benefits of all EBX® features available on datasets.

Indexing and search strategies

This feature is documented in the chapter [Search](#).

Business objects

A business object is an abstract layer that allows manipulating records from several tables within the same data model. The definition of business objects, named business object model, can be performed in the data model via the 'Extensions' section. A business object model is defined on a main table and contains child entities that are defined using the relationships from this main table.

Note:

Only the addon EBX® Match and merge is concerned by business objects. That is, business objects are not used in the core features of EBX®. This extension can be enabled only if the addon EBX® Match and merge is registered in the application server.

Create a business object model from the *Business objects > Definitions* section of the navigation pane, by clicking on the arrow



located to the left of *Business Objects*, then selecting the *Create Business Object* option. Follow the creation wizard to create a business object model. A business object model defines the following information:

Business object name	Business object model's name. The name must be unique in the context of the data model since it is used to uniquely identify the business object model. That is, it is not allowed to create several business objects with the same name.
Label and description	Internationalized labels and descriptions to be displayed to end users.
Main table	Defines the main table of this business object model. Specified path must be absolute (must start with "/"). The main table of this business object model cannot be used if it is already a part of another business object model. That is, the main table cannot be used if it is already defined as the main table or as a child entity of another business object model.

Defining the structure of a business object model

A business object model is composed of a main table and child entities. Child entities are defined using the relationships contained in the main table.

Add a child entity under a business object model by clicking on the arrow



located to the left of the existing business object, and by selecting the *Add child entity* option. Then, follow the creation wizard to create a child entity. The following information must be set:

Relationship to table	This field defines the relationship to the table to add as a child entity. The path to the

relationship must be absolute (must start with "/"). Only foreign key constraints and associations by inverse foreign key are supported. Also these relationships must comply with the following constraints:

Foreign key constraints cannot be a list or under a list. The field that defines the foreign key constraint must be declared as unique. That is, it must either define a uniqueness constraint with a blocking validation mode or be the only primary key field of its container table. Associations by inverse foreign key must refer to a foreign key field that cannot be a list or under a list. Also, the target table of the relationship cannot refer to a table that is already used by this business object model or another one. That is, the target table cannot be used if it is already defined as the main table or as a child entity of a business object model.

Toolbars

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 'Extensions' section.

Add a toolbar from the *Toolbars* section of the navigation pane, by clicking on the arrow



located to the left of , then selecting the *Create toolbar* option. Follow the creation wizard to create a toolbar. A toolbar defines the following information:

Name	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.

Label and description	Internationalized labels and descriptions to be displayed to end users.
Default template	Allows to create a toolbar with the structure of a default toolbar.
Locations	Specifies the locations where the toolbar can be used in associated datasets. Tabular view top: Defines the toolbar to use on top of the default table view. Tabular view row: Defines the toolbar to use on each row of the default table view. Tile view item: Defines the toolbar to use on each item of the tile view. Record top: Defines the toolbar to use in the record form. Hierarchy top: Defines the toolbar to use in the default hierarchy view of the table. Known limitations A toolbar defined to be used as "Tile view item" cannot have more than six items and it can only define action items. The display of those action buttons are always in mode "Icon only".

Defining the structure of a toolbar

A toolbar can define the following elements:

- [Action button](#)
- [Menu button](#)
- [Separator](#)
- [Menu group](#)

- Action menu item
- Sub menu item

Add one of these elements under a toolbar or to an existing element by clicking on the arrow



located to the left 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:

Service	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.
Label and description	Internationalized labels and descriptions to be displayed to end users.
Layout	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.
Icon	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.
Relief	Defines how this button will display. The button can be displayed as embossed or flat.

Is highlighted	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:

Label and description	Internationalized labels and descriptions to be displayed to end users.
Layout	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.
Icon	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.
Relief	Defines how this button will display. The button can be displayed as embossed or flat.
Is highlighted	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:

Label and description	Internationalized labels and descriptions to be displayed to end users.
Group type	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.
Service group name	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.
Menu builder name	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).
Excluded services	Indicates the services to exclude from the group of reused services. These services will not be displayed to end users in associated datasets.
Excluded service groups	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.
Filtering policy	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:

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

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 left 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 left 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 left 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'** menu in the navigation pane.

See also:

- [Data model import and export](#)

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'** menu in the navigation pane.

See also:

- [Data model import and export](#)

See also:

- [Use of toolbars](#)

Entity mappings

You can refer to tables in SQL requests and data service operations using unique names instead of their paths. To accomplish this, you must define entity mappings to create an association between a table path and a name. When referencing a table you then use this name as a WSDL suffix in data service operations and an alias in SQL requests.

Note:

EBX® automatically creates an entity name for the tables defined in a data model. The last step of the table's path is used as the default entity name. Since entity names must be unique in a data model, an error is now raised during compilation if the model contains tables with the same name. To prevent an error in this case, you must explicitly set the entity names.

To define an entity mapping through the user interface, in the navigation pane when viewing a data model, navigate to *Extensions > Entity mappings* and create a new record. You can define the following properties:

Table path	Specifies the path of the table in the current data model that is to be referred by this name in SQL requests or WSDL operations.
Entity name	This name is used to suffix all WSDL operation names of the concerned table and as an alias in SQL requests. 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](#)

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 tables in the relational database. These 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 extensions' section 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:

Name	Name of the replication unit. This name identifies a replication unit in the current data model. It must be unique.
Dataspace	Specifies the dataspace relevant to this replication unit. It cannot be a snapshot.
Dataset	Specifies the dataset relevant to this replication unit.
Refresh policy	Specifies the data synchronization policy. The possible policies are: • On commit: 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. • On demand: The replicated table in the database is only updated when an explicit refresh operation is performed.
Tables	Specifies the tables in the data model to be replicated in the database. Table path: Specifies the path of the table in the current data model that is to be replicated to the database. Table name in database: 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.

Aggregated lists

Specifies the properties of the aggregated lists in the table that are replicated in the database.

Path: Specifies the path of the aggregated list in the table that is to be replicated to the database.

Table name in database: 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.

See also:

- [Replication](#)

Node confidentiality

The confidentiality of a node defines whether or not it can be used in a Query or in a Request when hidden for a specific user. By default, all nodes are confidential.

This can be configured in the DMA, by using the "Extensions" tab in the node form editor. It is possible to define the confidentiality at several levels:

- at a data model level (root node): defines the default confidentiality for the data model terminal nodes, unless a more specific definition is set (that is on the table or the node itself).
- at a table level (table node): defines the default confidentiality for the table terminal nodes, unless a more specific definition is set (that is on the node itself).
- at a terminal node level: defines the confidentiality of the node.

For more information, see [Defining confidentiality](#)

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'** menu in the workspace.

- [Validating a data model](#)
- [Data model import and export](#)
- [Duplicating a data model](#)
- [Deleting a data model](#)

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](#) for detailed information on incremental data validation.

Data model import and export

TIBCO EBX® includes built-in data model services to import from and export to XML Schema Document (XSD) files or to archive files (zip). Imports and exports can be performed from the **data model 'Actions'** menu of the target data model in the navigation pane. An 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 data model (XSD or archive). Similarly, during exports, the entire data model is included in the XSD or archive file.

When importing a data model, the corresponding XSD 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 data model. See [Data model properties](#) for more information on declaring modules.

To perform an import select 'Import data model' from the **data model 'Actions'** menu of the data model into which you are importing.

You can import an XML Schema Document (XSD) or an archive (zip file) 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 or archive file to import.

Note:

When importing an archive it must contain only one XML Schema Document (XSD) at its root. XML documents related to data model extensions must also be located at the root of the archive. These XML documents are automatically imported if they comply with the naming of the extensions supported and registered in EBX®.

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:

Module	Module in which the data model is packaged.
Module path	Path to the module containing the data model.
Source path	Path to Java source used to configure business objects and rules. This property is required if the data model being imported defines programmatic elements.
Model	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'.

To perform an export select 'Export data model' from the **data model 'Actions'** menu of the data model you want to export.

Duplicating a data model

To duplicate a data model, select 'Duplicate' from the **data model 'Actions'** menu for that data model. You must give the new data model a name that is unique in the repository.

Deleting a data model

To delete a data model, select 'Delete' from the **data model 'Actions'** 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](#) for more information on the publication process.

Publishing and versioning data models

- [Publishing a data model](#)
- [Versioning an embedded data model](#)

Publishing a data model

- [About publications](#)
- [Publication modes](#)
- [Embedded publication mode](#)
 - [Viewing and creating publications](#)

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](#) 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.

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.

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](#) 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](#) 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'** 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'** 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.

Versioning an embedded data model

- [About versions](#)
- [Accessing versions](#)
- [Working with versions](#)
- [Known limitations on data model versioning](#)

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.

Accessing versions

To see the existing versions of your data model, select 'Manage versions' from the **data model** **'Actions'** 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.

Working with versions

In the workspace, using the down arrow



menu next to each version, you can perform the following actions:

Access data model version	Go to the corresponding version of the data model.
Create version	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.
Set as default version	Sets the selected version as the default version opened when users access the data model.
Export archive	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. See Archives directory for more information.
Import archive	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.

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

- [Introduction to dataspaces](#)
- [Creating a dataspace](#)
- [Working with existing dataspaces](#)
- [Snapshots](#)

Introduction to dataspaces

- [Overview](#)

- What is a dataspace?
- Basic concepts related to dataspaces
- Using the Dataspaces area user interface


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](#)
- [snapshot](#)
- [dataset](#)
- [dataspace merge](#)
- [reference dataspace](#)

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 dataspace and lists its snapshots.

D-S 593	
Identifier	Tutorial
Type	Dataspace
Creation	on 07/10/2013 at 12:32:34 by Doc. Adm
Status	Open
History purge	Marked for history purge by Doc. Admini
Owner	Doc. Administrator (admin)
Loading strategy	On-demand loading and unloading
Child merge policy	Allows validation errors in result
Child dataspace sort policy	By creation date (asc)

See also:

- [Creating a dataspace](#)
- [Snapshots](#)

Related concepts:

- [Datasets](#)

Creating a dataspace

- [Overview](#)
- [Properties](#)

Overview

To create a new dataspace, 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 dataspaces in the repository, dataspaces are always a child of another dataspace.

Properties

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

Identifier	Unique identifier for the dataspace.
Owner	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.
Label	Localized label and description associated with the dataspace.

Working with existing dataspaces

- [Dataspace information](#)
- [Dataspace permissions](#)
 - [General permissions](#)
 - [Allowable actions](#)
- [Merging a dataspace](#)
 - [Initiating a merge](#)
 - [Reviewing and accepting changes](#)
 - [Finalizing a merge](#)
 - [Abandoning a merge](#)
- [Comparing a dataspace](#)
- [Validating a dataspace](#)
- [Dataspace archives](#)
 - [Exporting](#)
 - [Importing](#)
- [Closing a dataspace](#)

Dataspace information

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

Documentation	Localized labels and descriptions associated with the dataspace.
Child merge policy	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. The available merge policies are: • Allows validation errors in result: Child dataspace can be merged regardless of the validation result. This is the default policy. • Pre-validating merge: A child dataspace can only be merged if the result would be valid.
Current Owner	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.
Child dataspace sort policy	Defines the display order of child dataspace in dataspace trees. If not defined, the policy of the parent dataspace is applied. Default is 'by label'.
Change owner	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.
Change permissions	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.

Dataspace permissions

General permissions

Dataspace id	The dataspace to which the permissions will apply.
Profile selection	The profile to which the rule applies.
Restriction policy	Whether these permissions restrict the permissions assigned to a given user through policies defined for other profiles. See Restriction policy.
Dataspace access	The global access permission on the dataspace. Read-only • Can see the dataspace and its snapshots, as well as child dataspace, according to their permissions. • Can see the contents of the dataspace depending on their permissions; cannot make modifications. Write • Can see the dataspace and its snapshots, as well as child dataspace, according to their permissions. • Can modify the contents of the dataspace depending on their permissions.

Hidden

- Cannot see the dataspace nor its snapshots directly.
- From a child dataspace, the current dataspace can be seen but not selected.
- Cannot access the contents of the dataspace.
- Cannot perform any actions on the dataspace.

Allowable actions

Users can be allowed to perform the following actions:

Create a child dataspace	Whether the profile can create child dataspaces.
Create a snapshot	Whether the profile can create snapshots from the dataspace.
Initiate merge	Whether the profile can merge the dataspace with its parent.
Export archive	Whether the profile can perform exports.
Import archive	Whether the profile can perform imports.
Close dataspace	Whether the profile can close the dataspace.
Close snapshot	Whether the profile can close snapshots of the dataspace.
Rights on services	Specifies the access permissions for services.

Permissions of child dataspace when created

Specifies the default access permissions for child dataspace that are created from the current dataspace.

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:

1. 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.
2. Changes that were made in the child dataspace since its creation are integrated into its parent dataspace.
3. The child dataspace is closed.
4. The parent dataspace is unlocked.

Initiating a merge

To merge a dataspace into its parent dataspace:

1. Select that dataspace in the navigation pane of the Dataspace area.
2. 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 "*< 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](#)

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 dataspaces will remain until you unlock them in the Dataspaces 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 dataspaces. Performing the unlock from the parent dataspace only unlocks the parent dataspace, thus you need to unlock the child dataspace separately.

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](#)

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.

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](#) in the Administration Guide for more information.

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

Name of the archive to create	The name of the exported archive.
Export policy	Required.

	The default export policy is 'The whole content of the dataspace', which exports all selected data to the archive. 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.
Datasets to export	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.

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](#)

Snapshots

- [Overview of snapshots](#)
- [Creating a snapshot](#)
- [Viewing snapshot contents](#)
- [Snapshot information](#)
- [Comparing a snapshot](#)
- [Validating a snapshot](#)
- [Export](#)
- [Closing a snapshot](#)

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](#)

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:

Identifier	Unique identifier for the snapshot.
Label	Localized labels and descriptions associated with the snapshot.

Viewing snapshot contents

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

Snapshot information

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

Documentation	Localized labels and descriptions associated with the snapshot.
Current Owner	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.
Change owner	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.

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](#)

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.

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](#) in the Administration Guide for more information.

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

Name of the archive to create	The name of the exported archive.
Datasets to export	The datasets to export from this snapshot. For each dataset, you can choose whether

to export its data values, permissions, and information.

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](#)

Datasets

- [Introduction to datasets](#)
- [Creating a dataset](#)
- [Viewing data](#)
- [Editing data](#)
- [File import and export services](#)
 - [XML import and export](#)
 - [CSV import and export](#)
- [Working with existing datasets](#)
- [Dataset inheritance](#)

Introduction to datasets

- [Overview](#)
 - [What is a dataset?](#)
 - [Basic concepts related to datasets](#)

- [Using the Data user interface](#)

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](#)
- [dataset](#)
- [record](#)
- [field](#)
- [primary key](#)
- [foreign key](#)
- [table \(in dataset\)](#)
- [group](#)

Using the Data user interface

Datasets can be created, accessed and modified in the **Data** area using the [Advanced perspective](#) 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](#)
- [Quick Search](#)
- [Working with records in the user interface](#)
- [Inheritance](#)

Related concepts:

- [Data model](#)
- [Dataspace](#)

Creating a dataset

- [Creating a root dataset](#)
- [Creating an inheriting child dataset](#)

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'**



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.

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.

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'**



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](#)

Viewing data

TIBCO EBX® offers different ways to list records. This chapter presents how to sort, search, and display records in varying ways, and according to different user profiles, through to the concept of 'Views'.

- ['View' menu](#)
- [Sorting data](#)
 - [Sorting by column](#)
- [Quick Search](#)
 - [Special characters](#)
- [Searching and filtering data](#)
 - [Search on a field](#)
 - [Validation filter](#)

- Custom table searches
- Views
 - View description
 - Simple tabular views
 - Hierarchical views
 - Tile views
 - View sharing
 - View publication
- Views management
 - Manage recommended views
 - 'View' menu toolbar
- Grid edit
 - Copy/paste
- History

'View' menu

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

Views are managed directly in the 'View' menu toolbar, available on each listed view: ['View' menu toolbar](#).

Views can also be grouped. An administrator has to beforehand define groups 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'](#) for more information.

Sorting data

Sort criteria control the order in which records are presented.

Use the 'Select and Sort' button at the top left of the table to define specific sorting criteria.

There are two types of sorting:

- sorting by relevance, used in the [quick search](#),

- sorting by column.

Sorting by column

The 'Sorting criteria' dialog box offers:


- on the left, the list of sorted columns,
- on the right, the list of unsorted columns.

Use the ← → arrow buttons to toggle columns from one list to another.

Use the ↑ ↓ arrow buttons to change the priority order of the column.

To change the sort order of a column, use the 'ASC' (ascending) or 'DESC' (descending) button that appears on the mouse-over.

Quick Search



The quick search is used to easily find a result in a tabular, a hierarchical, or a tile view.

It does not differentiate between upper and lower case. It allows you to search for several terms at once (separated by spaces). By default, the records found are sorted by relevance.

See also:

- [Search](#)

Special characters

The quick search offers special characters to refine your search.

+ ... <i>required</i>	Makes it required that the word be present in the result. Prevents the search engine from excluding this word from the search. Note: Works with +" phrase " or +(group). <i>Example:</i> +flute +bach ↳ Finds results with required <i>flute</i> and required <i>bach</i>. Results with only <i>flute</i> or only <i>bach</i> will be ignored.
- ... <i>not</i>	Excludes the word from the result. Note: Works with -" phrase " or -(group). <i>Example:</i> bach -flute ↳ Find results with <i>bach</i>, but without <i>flute</i>.
... ~ <i>fuzzy</i>	Specifies that the search can change 2 characters of the word to find it. To allow for a single character to change, use ~1. Note: Works with " phrase "~, to change 2 words in the phrase. <i>Example:</i> handel~ ↳ Find results with a word that is <i>handel</i>, with 1 or 2 letters differing.

? <i>single joker</i>	Replaces an unknown character. <i>Example:</i> ?uttle ↳ Finds results with a word starting with any character, and ending with <i>uttle</i>.
* <i>extended joker</i>	Replaces several unknown characters. <i>Example:</i> rachmanino* ↳ Find results with a word beginning with <i>rachmanino</i>.
"..." <i>phrase</i>	Find the exact match of the phrase. Note: Can be surrounded by + - ~. <i>Example:</i> "Johann Sebastian Bach" ↳ Find results containing exactly <i>Johann Sebastian Bach</i>.
(...) <i>group</i>	Allows grouping words to apply a special character + or - to them. Note: You can make groups of groups. <i>Example:</i> bach +(flute piano) ↳ Find the results with possibly <i>bach</i>, and necessarily <i>flute</i> or <i>piano</i>.

Note:

These special characters can also be used in the documentation search engine.

Searching and filtering data

The search pane is hidden by default and accessible via the



(/content/dam/onebx/us/en/ebx/webhelp/6-1-0/resources/images/symbols/Open Search Pane.svg) icon located to the right of the quick search in the toolbar of the tabular, the hierarchical or the tile view.

The quick search and the criteria lines combine to narrow the search (restricting the result to fewer and fewer records).

It is possible to deactivate a criteria line by unchecking it. The deactivated criteria are not kept during a save.

The trashcan button



at the end of the line of each criterion permanently deletes the criterion.

To save the filter applied to a search, use the 'Save' button. Saving takes into account the quick search and all active criteria.

To recall a saved filter, use the 'Load' button. Loading replaces the quick search and the whole criteria panel. Click on the 'Apply' button to start the new search.

When a view is applied, it ensures that it is displayed according to its configuration. All existing criteria in the search panel are therefore removed. The view can contain a set of search criteria, which are applied together with the at the same time as the view.

Some operators (such as 'matches') allow to use Lucene regular expressions. See [technical specifications of Lucene's regex pattern](#) for more information.

Search on a field

All searchable fields are available.

Validation filter

In field selection, the validation criteria display the records as of the last validation performed.

Note:

This filtering 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.

To filter on the validation severity level (independent from validation Message), use the 'Severity' validation criterion. Available levels are: 'Errors', 'Warnings' and 'Information'.

To filter on the validation message (independent from validation Severity level), use the 'Message' validation criterion.

Custom table searches

For backward compatibility, the feature for custom searching and filtering records is still operational and accessible via the icon



in the workspace. The icon and feature are only available when at least one custom filter exists.

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

See also:

- [Customizing table filter](#)

Views

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

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*.

Three 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.

- 'Tile view': A view which displays the records in the form of cards that are positioned in rows and columns. This view allows to show information in an eye catching form, sort, and filter the records displayed.

View description

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

Documentation	Localized label and description associated with the view.
Owner	Name of the owner of the view. This user can manage and modify it. (Only available for administrators and dataset owners)
Share with	Other profiles allowed to use this view from the 'View' menu. Note: Requires a permission, see Views permissions .
View mode	Simple tabular view, hierarchical view or tile view.
View group	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.

Displayed columns	Specifies the columns that will be displayed in the table.
Sorted columns	Specifies the sort order of records in the table. See Sorting data .

Filter	Defines filters for the records to be displayed in this table view. See Criteria editor.
Initial request	When set to manual, users of this view will have to apply a filter or a sort criteria before they can view data.
Pagination limit	Forces a limit to the number of visible records.
Grid edit	If enabled, users of this view can switch to grid edit, so that they can edit records directly from the tabular view.
Disable create and duplicate	If yes, users of this view cannot create nor duplicate records from the grid edit.
Toolbar on top of table	Defines the toolbar to use on top of this table view.
Toolbar on table row	Defines the toolbar to use on each row of this table view.

Hierarchical views

A hierarchy view displays your data in a tree-like representation. Use hierarchy views to visualize different relationships between data, or to group common values. To create a hierarchy view, you define a dimension that determines what data will be included and how it displays. Consider the following example and image where the dimension is Category --> Brand --> Product and displays Products by Brand by Category.




Hierarchy dimensions and members

Each hierarchy is defined by its dimension. The dimension is subdivided into dimension levels. Members are hierarchy nodes and represent a record, or a field value. A member can be a parent of other members, or a leaf node with no children. The table from which you create the hierarchy view is called the target table. Members of the target table are lowest-level leaves, except in the case of a recursive hierarchy.

In the following image:

- Grocery and Household are members of the Category dimension level. This is the top dimension level.
- ACME is a member of the Brand dimension level and is at an intermediate level. A dimension can have many intermediate levels.
- The individual products listed at the Product dimension level are leaf members from the target table. This is the bottom dimension level.



Creating a hierarchy view

This topic provides high-level instructions for creating a hierarchy view. See the following topics for additional details about creating views of specific hierarchy types:

- Ragged hierarchies
- Group by hierarchies
- Time-based hierarchy dimensions

To create a hierarchy view:

1. From the target table's **Actions** menu, select **Create a new view**.
2. After specifying the following information, click **Next**:
 - Add any view-related documentation such as a label and description.
 - Set access permissions.
 - For the **View mode** property, select **Hierarchical view**.
3. Use the **Dimension** box to define the hierarchy by working backwards through the dimension levels:
 - a. Bottom level: The members of this level are pre-determined based on the target table.
 - b. Intermediate levels: Start with the attribute that corresponds to the next highest dimension level and continue to expand until you navigate to the highest level.
 - c. Top level: Select the attribute corresponding to the top dimension level.
4. Review the required options and after making desired changes, click **Next**.
5. Specify labels, filters, and sort options for the members of each dimension level:
 - Navigate between dimension levels by selecting them in the breadcrumb at the top of the screen.
 - Make adjustments to labels, filters, and sort strategies as desired.
6. Click **Save and quit** to display the hierarchy view.

Ragged hierarchies

A ragged hierarchy includes members where their parent skips one level. An example of this type of hierarchy is shown by using the Country, State, and City dimension. Some countries do not have states, but they all have cities. In the case of the latter, the city members directly roll up to the Country dimension level.

Attention:

A ragged hierarchy only supports skipping one level.

To display ragged hierarchies in EBX®:

- In the data model, you must define all foreign key relationships from the target table. Using the example above, the City table must hold a foreign key to the State and Country tables. Of course, for cities with no state, this field is left undefined.
- When configuring your hierarchy view, enable the **Allow skipping level** option under **Advanced properties**. This property is highlighted in the image below:


The screenshot shows the configuration interface for a hierarchy view. The 'Dimension' section is at the top, showing a tree structure with 'Store' as the root, followed by 'CityFK (City)', 'CityName', 'Creator', 'Last updater', and 'CountryFK (Country)'. Below this, there are several configuration options:

- Display records in a new window: ☐ Yes ☒ No
- Prune hierarchy: ☐ Yes ☒ No
- Display orphans: ☒ Yes ☐ No
- Display root node: ☒ Yes ☐ No
- Root node label: English (United States)
- Toolbar on top of hierarchy: [not defined]
- Advanced properties:
 - Remove recursive root leaves: ☐ Yes ☒ No
 - Detect cycle: ☐ Yes ☒ No
 - Allow skipping level: ☒ Yes ☐ No
 - Detect leaf: ☒ Yes ☐ No
 - Allow children count: ☐ Yes ☒ No

Group by hierarchies

The dimension in a Group By hierarchy is defined using columns located in the same table. Allowed column data types are foreign keys or the following enumerable types: String, Integer, Enumeration, and Boolean. This provides a powerful tool for defining hierarchy views that group common values together. A group by hierarchy also provides flexibility when setting the order for

dimension levels. For example, a Product table might include the following columns: Brand, Segment, Type, and Packaging. Use these columns in any order to show products by: Segment --> Type --> Packaging, or Brand --> Segment --> Type --> Packaging.




Time-based hierarchy dimensions

Use fields with a Date data type to define a time-based hierarchy dimension. When configuring the view, choose from several predefined display options. Some options subdivide the time into smaller units. For example, you can specify that a project's Due Date field is used to display all project deadlines by month in Q1 of 2022.

To use a time-based dimension in a hierarchy view:

1. Create a new hierarchy view configuration:
 - a. When viewing a table, open the **View** menu and select **Create a new view**.
 - b. Provide a name and optionally specify share options.
 - c. For the **View mode**, select **Hierarchical view** and click **Next**.
2. Ensure the field used for the hierarchy's root dimension is of the Date data type and click **Next**.

- Click on the Date field at the top of the screen in the hierarchy structure breadcrumb and use the **Time dimension** menu to specify how you want the dimension to display.



- Save and view the hierarchy.

Hierarchical view configuration options

This form allows configuring the hierarchical view options.

Display records in a new window	If 'Yes', a new window will be opened with the record. Otherwise, it will be displayed in a new page of the same window.
Prune hierarchy	If 'Yes', hierarchy nodes that have no children and do not belong to the target table will not be displayed.
Display orphans	If 'Yes', hierarchy nodes without a parent will be displayed.
Display root node	If 'No', the root node of the hierarchy will not be displayed in the view.
Root node label	Localized label of the hierarchy root node.
Toolbar on top of hierarchy	Allows setting the toolbar on top of the hierarchy.

Detect cycle	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.
Detect leaf	Allows detecting whether the member is a leaf or not. The leaf detection is very costly for large volumes of data. Thus, it is recommended to disable this option when the query response is delayed to display the hierarchy view. This property is always disabled for orphans' parent members.

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](#)

Sort strategy

For each dimension level, it is possible to choose one of the following sort strategies:

Default	Nodes are sorted by label in alphabetical order
Sort by columns	Nodes are sorted by selected column(s). The direction (ascending/descending) can be chosen for each column.

Sort by ordering field

Nodes are sorted by a hidden numeric field, which allows users to dynamically change the order of sibling nodes in the hierarchy view. This strategy is available only if there is at least one 'Hidden' numeric field in the table.

In order to enable this option, 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.

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.

Tile views

This view displays records as tiles arranged in a grid. Each tile can include the following information: Title (required), Subtitle, Image, and description. For the titles and description you can: select fields to automatically populate the values, manually specify the text that displays, or use a combination of both methods. Additionally, you can define sort and filter criteria that applies to the view. Toolbars can also be added. The following table lists and defines the properties available when creating a tile view:

Title	Defines the tile title displayed on first line.
Subtitle	Defines the subtitle displayed on second line.
Image	Defines the tile image.
Image > Default image URL	Defines the fallback image URL.
Image > URL	Defines the image URL.
Image > Height (px)	Sets the height of the image slot in pixels (between 10 and 250). The slot width will always be 180 pixels. If the dimensions of the image exceed the size of the intended slot, it will be adjusted according to the cropping rules defined by the view.
Image > Cropping	Defines the cropping of the image. If "Cover", the image will be scaled up or down so that occupy all the available space. In this case, the image might be truncated. If "Scale down", the image will be scaled down only if it overflows the available space. Whatever the value, the image will always be centered.
Tile description	Defines the tile text.
Tile description > Number of displayed lines	Defines the number of lines (between 3 and 6).

Tile description > Line wrap	If "Yes", each row is displayed entirely (with line wrap) up to the defined limit of the number of lines. The excess is then truncated. If "No", each line is truncated.
Tile description > Content	Defines the tile text content.
Sort criteria	Specifies the sort order of records in the view. See Sorting data .
Filter	Defines filters for the records to be displayed in this tile view. See Criteria editor .
Toolbar on top of table	Defines the toolbar to use on top of this table view.
Tile toolbar	Defines the toolbar to use on each item of this tile view.

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 the 'View menu', click on the Edit button displayed on the mouseover of a listed view and add a 'Publication name' to the view.

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.

'View' menu toolbar

The 'View' menu toolbar offers the following actions:

Edit	Click on the 'Edit' button of the targeted view's toolbar to access the editable form.
Duplicate	Click on the 'Duplicate' button of the targeted view's toolbar to duplicate the view. The new view creation form pre-populates the field values from the view being duplicated.
Delete	Click on the 'Delete' button of the targeted view's toolbar to delete the view.
Define this view as my favorite	Click on the 'Define this view as my favorite' button of the targeted view's toolbar. The favorite view will automatically be applied when accessing the table. Click a second time on the button to remove the view as the user's favorite view.

Grid edit

The grid edit feature allows modifying 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](#)

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.

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](#) 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:

History in current dataspace

The table history view displays operations on the current branch. This is the default mode.

History in current dataspace and ancestors	The table history view displays operations on the current branch and on all its ancestors.
History in current dataspace and merged children	The table history view displays operations on the current branch and on all its merged children.
History in all dataspaces	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](#)

Editing data

- [Working with records in the user interface](#)
 - [Creating a record](#)
 - [Updating an existing record](#)
 - [Duplicating a record](#)
 - [Deleting records](#)
 - [Comparing two records](#)
- [Importing and exporting data](#)
- [Restore from history](#)

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 or a tile 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.

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](#)
- [XML Services](#)

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.isHistoryRestore
```

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](#).
- 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](#)

File import and export services

- [XML import and export](#)
- [CSV import and export](#)

XML import and export

- [Introduction](#)
- [Imports](#)
 - [Import mode](#)
 - [Insert and update operations](#)
 - [Optimistic locking](#)
- [Exports](#)
- [Handling of field values](#)
 - [Date, time & dateTime format](#)
- [Known limitations](#)
 - [Association fields](#)
 - [Selection nodes](#)

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*.

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.

Insert mode	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.
Update mode	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.
Update or insert mode	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.
Replace (synchronization) mode	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](#) and [insert](#), as well as `ImportSpec.setByDelta` in the Java API for more information.

STATE IN SOURCE XML DOCUMENT	BEHAVIOR
Element does not exist in the source document	If 'by delta' mode is disabled (default): Target field value is set to one of the following: • If the element defines a default value, the target field value is set to that default value. • If the element is of a type other than a string or list, the target field value is set to <code>null</code>. • If the element is an aggregated list, the target field value is set to an empty list. • If the element is a string that distinguishes <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 distinguish between the two, an empty string.

STATE IN SOURCE XML DOCUMENT	BEHAVIOR
	- If the element (simple or complex) is hidden in data services, the target value is not changed. See also: Hiding a field in Data Services Note: The user performing the import must have the permissions necessary to create or change the target field value. Otherwise, the value will remain unchanged. If 'by delta' mode has been enabled through data services or the Java API: - For the <code>update</code> operation, the field value remains unchanged. - For the <code>insert</code> operation, the behavior is the same as when <code>byDelta</code> mode is disabled.
Element exists but is empty (for example, <code><fieldA/></code>)	- For nodes of type <code>xs:string</code> (or one of its sub-types), the target field's value is set to <code>null</code> if it distinguishes <code>null</code> from an empty string. Otherwise, the value is set to empty string. - For non-<code>xs:string</code> type nodes, an exception is thrown in conformance with XML Schema. See also: TIBCO EBX® whitespace management for data types

STATE IN SOURCE XML DOCUMENT	BEHAVIOR
Element is present and <code>null</code> (for example, <code><fieldA xsi:nil="true"/></code>)	The target field is always set to <code>null</code> except for lists, for which it is not supported. 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> .

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.

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:

Download file name	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.
User-friendly mode	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. Note: If this option is selected, the exported file will not be able to be re-imported.
Include technical data	Specifies whether internal technical data will be included in the export. Note: If this option is selected, the exported file will not be able to be re-imported.
Is indented	Specifies whether the file should be indented to improve its readability by a human.
Omit XML comment	Specifies whether the generated XML comment that describes the location of data and the date of the export should be omitted from the file.

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

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.

CSV import and export

- [Introduction](#)
- [Exports](#)
- [Imports](#)
- [Handling of field values](#)
 - [Aggregated lists](#)
 - [Hidden fields](#)
 - ['Null' value for strings](#)
 - [Date, time & dateTime format](#)
- [Known limitations](#)
 - [Aggregated lists of groups](#)

- Terminal groups
- Column label headers
- Association fields
- Selection nodes

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](#)

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:

Download file name	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.
File encoding	Specifies the character encoding to use for the exported file. The default is UTF-8.

Enable inheritance	In order to consider the inheritance during a CSV export, the option has to be defined in the model. For more information on inheritance, see Inheritance and value resolution. Specifies if inheritance will be taken into account during a CSV export. 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. If inheritance is disabled, resolved values of fields are exported and occulted records are ignored. By default, this option is disabled. Note: Inheritance is always ignored, if the table dataset has no parent or if the table has no inherited field.
User-friendly mode	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. Note: If this option is selected, the exported file will not be able to be re-imported.
Include technical data	Specifies whether internal technical data will be included in the export.

	Note: If this option is selected, the exported file will not be able to be re-imported.
Column header	Specifies whether or not to include column headers in the CSV file. • No header • Label: 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. • 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 <i>Administration > User interface</i>.
List separator	Specifies the separator to use for values lists. The default separator is line return, it can be modified under <i>Administration > User interface</i>.

Programmatic CSV exports are performed using the classes `ExportSpec` and `ExportImportCSVSpec` in the Java API.

Imports

Download file name	Specifies the name of the CSV file to be imported.

Import mode	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. • Insert mode: 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. • Update mode: 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. • Update or insert mode: 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. • Replace (synchronization) mode: 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.
File encoding	Specifies the character encoding to use for the exported file. The default is UTF-8.
Enable inheritance	In order to consider the inheritance during a CSV import, the option has to be defined in the model.

	For more information on inheritance, see Inheritance and value resolution and <code>ExportImportCSVSpec.setInheritanceEnabled</code>. 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. Note: Inheritance is always ignored if the dataset of the table has no parent or if the table has no inherited field.
Column header	Specifies whether or not to include column headers in the CSV file. • No header • Label: 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. • 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 <i>Administration > User interface</i>.

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 imports are performed using the classes `ImportSpec` and `ExportImportCSVSpec` in the Java API.

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.setNullStringEncoded` 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

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](#).

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.

Working with existing datasets

- [Validating a dataset](#)
- [Duplicating a dataset](#)
- [Deactivating a dataset](#)
- [Managing dataset permissions](#)

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](#) for detailed information about incremental data validation.

Duplicating a dataset

To duplicate an existing dataset, select it from the '[Select dataset](#)'



menu in the navigation pane, then select **Actions > Duplicate**.

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**.

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](#)

Profile	Defines the profile to which these permissions apply.
Restriction policy	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. See Resolving user-defined rules.
Dataset actions	Specifies the permissions for actions on the dataset.
Create a child dataset	Indicates whether the profile can create a child dataset. Inheritance also must be activated in the data model.
Duplicate the dataset	Indicates whether the profile can duplicate the dataset.
Delete the dataset	Indicates whether the profile can delete the dataset.
Activate/deactivate the dataset	Indicates whether the profile can modify the <i>Activated</i> property in the dataset information. See Deactivating a dataset .

Create a view	Indicates whether the profile can create views and hierarchies in the dataset.
Tables policy	Specifies the default permissions for all tables. Specific permissions can also be defined for a table by clicking the '+' button.
Create a new record	Indicates whether the profile can create records in the table.
Overwrite inherited record	Indicates whether the profile can override inherited records in the table. This permission is useful when using dataset inheritance.
Occult inherited record	Indicates whether the profile can occult inherited records in the table. This permission is useful when using dataset inheritance.
Delete a record	Indicates whether the profile can delete records in the table.
Values access policy	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.
Rights on services	This section specifies the access permissions for services. A service is not accessible to a profile if it is crossed-out.

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](#)
- [Dataset inheritance structure](#)
- [Value inheritance](#)
 - [Record inheritance](#)

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.

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.

Root	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.
Inherited	An inherited record is one that is defined in an ancestor dataset of the current dataset.
Overwritten	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.

Occulted	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
Root	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.
Inherited	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.
Overwritten	Not applicable. Cannot create a new record if the primary key is already used	An overridden record can be returned to the inherited state, but its modified value will be lost. Individual	Deleting an overwritten record changes its state to occulted.

STATE	CREATE	MODIFY VALUE	DELETE
	in the current dataset.	values in an overridden record can be set to inheriting or can be modified.	
Occulted	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](#)

Collaborative features (Team Up)

- [Introduction to collaborative features \(Team Up\)](#)
- [Comments](#)
- [Ratings](#)
- [Tags](#)

Introduction to collaborative features (Team Up)

- [Overview](#)
- [Accessing the Team Up features](#)

Overview

Use Team Up to collaborate and engage with others. Leave comments on records, rate records, and assign tags to records. These interactions can promote a better overall understanding of data.

The following provides an overview of Team Up features:

- **Comments:** Leave comments on records, reply, and react to other's comments. Assign moderators to tables that can manage comments from other users.
- **Tags:** Apply tags to records. Tags are classified in tag catalogs. Assign an administrator to manage tag catalogs.
- **Ratings:** Leave a rating from 1 to 5 stars on each record. When leaving a rating, optionally add a note to explain it.

Accessing the Team Up features

To access the Team Up features available for a record:

1. Open a record details page.
2. In the right sidebar, click the icon corresponding to the Team Up feature you want to use.

The screenshot displays the 'Authors' record details page in TIBCO EBX 6.1.0. The interface includes a top navigation bar with various icons and a left sidebar with a menu. The main content area shows the 'Authors' section with fields for Id, Lastname, Firstname, Address, Birth date, Death date, City, and Country. Below these fields is a 'Books' section with a table showing no records found. The right sidebar contains three icons: a speech bubble (Comments), a star (Rating), and a tag (Tags). The 'Authors' section is active, showing fields for Id, Lastname, Firstname, Address, Birth date, Death date, City, and Country. Below these fields is a 'Books' section with a table showing no records found.

See also:

- Comments
- Ratings
- Tags

Comments

- Overview
- Adding a comment
- Editing a comment
- Deleting a comment
- Replying to a comment
- Reacting to a comment
- Moderating a comment

Overview

Use the Team Up comments feature to discuss the details of a record. Add your own comments, or reply and react to comments from other users.

The screenshot displays the TIBCO EBX 6.1.0 user interface. On the left, a sidebar contains a 'Tutorial' section and a list of actions: 'Titles', 'Publishers', 'Authors' (selected), 'Royalties', 'Contracts', and 'BookRental'. The main content area is titled 'Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]'. It features a form for editing author details, including fields for 'Id' (00-0-001), 'Lastname' (lastNameAuthor1), 'Firstname' (firstNameAuthor1), 'Address' (addressAuthor1), 'Birth date', 'Death date', 'City' (cityAuthor1), and 'Country' (France). Below the form is a table for 'Books' with columns 'Id', 'Title', 'Type', 'Creation date', and 'Publisher'. The table currently shows 'No records found.' At the bottom of the form are buttons for 'Save', 'Save and close', 'Revert', and 'Close'. On the right, a 'Comments' sidebar is open, showing a list of comments. The first comment is from 'User1 User1' on 10/19, 17:54, with the text 'Lorem ipsum dolor sit amet.' and 1 like, 0 dislikes, and 0 replies. The second comment is from 'User2 User2' on 10/19, 18:07, with the text 'Consectetur adipiscing elit. (moderated)' and 0 likes, 0 dislikes, and 1 reply. The third comment is from 'User3 User3' on 10/19, 18:15, with the text 'Nam et mi eu leo mollis aliquet. Duis at q uam.' and 0 likes, 1 dislike, and 0 replies. At the bottom of the comments sidebar is a text input field labeled 'Write your comment' and a 'Post' button.

Adding a comment

Add a comment using the text input at the bottom of the **Comments** sidebar.

The screenshot displays the TIBCO EBX 6.1.0 user interface. On the left, the 'Master Data - Reference' sidebar shows a list of entities: Tutorial, Actions, Titles, Publishers, Authors (selected), Royalties, Contracts, and BookRental. The main area shows the 'Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]' form. The 'Authors' tab is active, showing fields for Id (00-0-001), Lastname (lastNameAuthor1), Firstname (firstNameAuthor1), Address (addressAuthor1), Birth date, Death date, City (cityAuthor1), and Country (France). Below these fields is a 'Books' table with columns: Id, Title, Type, Creation date, and Publisher. The table shows 'No records found.' At the bottom of the form are buttons: Save, Save and close, Revert, and Close. On the right, the 'Comments' sidebar is open, showing a list of comments. The bottom of the sidebar features a text input field labeled 'Write your comment' with a red box around it, and a 'Send' button.

Editing a comment

Edit your comments by clicking the editing button in the toolbar (highlighted below). Note that comments edited by a moderator cannot be edited.

For details about comment moderation, see [Moderating a comment](#).

The screenshot shows the TIBCO EBX 6.1.0 interface. On the left, the 'Master Data - Reference' menu is open, and 'Authors' is selected. The main area displays the 'Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]' form. The 'Comments' panel on the right shows two comments. The first comment, by 'User1 User1', has a toolbar with a delete icon (trash can) highlighted by a red box. The second comment, by 'User3 User3', is also visible.

Deleting a comment

Delete one of your comments by clicking the delete button in the toolbar located under the comment.

The screenshot shows the TIBCO EBX 6.1.0 interface. On the left, the 'Master Data - Reference' menu is open, and 'Authors' is selected. The main area displays the 'Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]' form. The 'Comments' panel on the right shows two comments. The first comment, by 'User1 User1', has a toolbar with a delete icon (trash can) highlighted by a red box. The second comment, by 'User3 User3', is also visible.

Replying to a comment

Reply to a comment using the reply button in the toolbar located under the comment.

Master Data - Reference ▾ Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Tutorial ▾

Actions ▾

Titles

Publishers

▾ Authors

lastNameAuthor1 firstNameAuthor1 [00-0-001]

Royalties

Contracts

BookRental

Authors EditorWithContracts

Id 00-0-001

Lastname * lastNameAuthor1

Firstname * firstNameAuthor1

Address addressAuthor1

Birth date / /

Death date / /

City cityAuthor1

Country France

Books + Actions ▾ No records ▾

Id	Title	Type	Creation date	Publish
No records found.				

Save Save and close Revert Close

Comments

UU User1 User1 10/19, 17:54

Lorem ipsum dolor sit amet.

1 0 0

1 reply to display ▾

UU User3 User3 10/19, 18:15

Nam et mi eu leo mollis aliquet. Duis at q uam.

0 1 0

Write your comment

Reacting to a comment

React to comments using the available reaction buttons in the toolbar located under the comment. Change your reaction by clicking a different reaction; remove your reaction by clicking the same reaction again.

Master Data - Reference ▾ Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Tutorial ▾

Actions ▾

Titles

Publishers

▾ Authors

lastNameAuthor1 firstNameAuthor1 [00-0-001]

Royalties

Contracts

BookRental

Authors EditorWithContracts

Id 00-0-001

Lastname * lastNameAuthor1

Firstname * firstNameAuthor1

Address addressAuthor1

Birth date / /

Death date / /

City cityAuthor1

Country France

Books + Actions ▾ No records ▾

Id	Title	Type	Creation date	Publish
No records found.				

Save Save and close Revert Close

Comments

UU User1 User1 10/19, 17:54

Lorem ipsum dolor sit amet.

1 0 0

1 reply to display ▾

UU User3 User3 10/19, 18:15

Nam et mi eu leo mollis aliquet. Duis at q uam.

0 1 0

Write your comment

Moderating a comment

Moderate and edit other's comments by clicking the moderate icon highlighted below. When editing other's comments you can leave a note that indicates why the original content was moderated.

The screenshot displays the TIBCO EBX 6.1.0 user interface. On the left, a sidebar contains navigation links: Master Data - Reference, Tutorial, Actions, Titles, Publishers, **Authors** (selected), lastNameAuthor1 firstNameAuthor1 [00-0-001], Royalties, Contracts, and BookRental. The main area shows the 'Authors' record editor for 'lastNameAuthor1 firstNameAuthor1 [00-0-001]'. The editor has two tabs: 'Authors' and 'EditorWithContracts'. The 'Authors' tab is active, showing fields for Id (00-0-001), Lastname (lastNameAuthor1), Firstname (firstNameAuthor1), Address (addressAuthor1), Birth date, Death date, City (cityAuthor1), and Country (France). Below these fields is a 'Books' section with a table header (Id, Title, Type, Creation date, Publisher) and a message 'No records found.' At the bottom of the editor are buttons: Save, Save and close, Revert, and Close. On the right, a 'Comments' sidebar is open, showing two comments. The first comment is from 'User1 User1' (10/19, 17:54) with the text 'Lorem ipsum dolor sit amet.' and a red warning icon. The second comment is from 'User3 User3' (10/19, 18:15) with the text 'Nam et mi eu leo mollis aliquet. Duis at q uam.' and a warning icon. At the bottom of the comments sidebar is a text input field 'Write your comment' and a submit button.

For details about how to manage moderators, see [Managing user access and Team Up features](#).

See also:

- [Ratings](#)
- [Tags](#)

Ratings

- [Overview](#)
- [Rating a record and editing a rating](#)
- [Deleting a rating](#)

Overview

Give records ratings by leaving up to 5 stars and optionally leave an explanation for the rating. The **Ratings** sidebar displays:

- A record's overall rating.
- The individual ratings supplied by other users.
- The comments related to those ratings.

Master Data - Reference

Tutorial

Actions

Titles

Publishers

Authors

lastNameAuthor1 firstNameAuthor1

Royalties

Contracts

BookRental

Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Authors

EditorWithContracts

Id

00-0-001

Lastname

lastNameAuthor1

Firstname

firstNameAuthor1

Address

addressAuthor1

Birth date

/ /

Death date

/ /

City

cityAuthor1

Country

France

Books

+

Actions

No records

Id

Title

Type

Creation date

Publisher

No records found.

Save

Save and close

Revert

Close

Ratings

4.0

★

★

★

★

★

★

2 ratings

UU

User2 User2

★

★

★

★

★

a few minutes ago

Lorem ipsum dolor sit amet.

UU

User1 User1

★

★

★

★

★

10/25, 10:46

Your rating

★

★

★

★

★

Explain your rating

Rating a record and editing a rating

Use the inputs at the bottom of the **Rating** sidebar to add your rating and optionally leave a comment. Perform the same actions to edit your rating.

Master Data - Reference

Tutorial

Actions

Titles

Publishers

Authors

lastNameAuthor1 firstNameAuthor1

Royalties

Contracts

BookRental

Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Authors

EditorWithContracts

Id

00-0-001

Lastname

lastNameAuthor1

Firstname

firstNameAuthor1

Address

addressAuthor1

Birth date

/ /

Death date

/ /

City

cityAuthor1

Country

France

Books

+

Actions

No records

Id

Title

Type

Creation date

Publisher

No records found.

Save

Save and close

Revert

Close

Ratings

4.0

★

★

★

★

★

★

2 ratings

UU

User2 User2

★

★

★

★

★

yesterday

Lorem ipsum dolor sit amet.

UU

User1 User1

★

★

★

★

★

10/25, 10:46

Your rating

★

★

★

★

★

Explain your rating

164

Deleting a rating

Delete your rating using the delete icon located to the right of your rating. Note that your rating always displays at the top of the individual ratings list.

The screenshot shows the TIBCO EBX interface with the 'Authors' sidebar on the left. The main area displays the 'Author' record editor for 'lastNameAuthor1 firstNameAuthor1 [00-0-001]'. The 'Ratings' panel on the right shows a 4.0 average rating with 2 ratings. A star icon is highlighted in the 'Ratings' panel, indicating the delete action.

See also:

- [Comments](#)
- [Tags](#)

Tags

- [Overview](#)
- [Applying tags](#)
 - [Applying existing tags](#)
 - [Applying a new tag](#)
- [Removing a tag](#)

Overview

Use the **Tags** sidebar to assign one or more tags to a record.

Master Data - Reference ▾ Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Tutorial ▾

Actions ▾

Titles

Publishers

▼ Authors

lastNameAuthor1 firstNameAuthor1

Royalties

Contracts

BookRental

Authors EditorWithContracts

Id 00-0-001

Lastname lastNameAuthor1

Firstname firstNameAuthor1

Address addressAuthor1

Birth date / /

Death date / /

City cityAuthor1

Country France

Books + Actions No records

Id	Title	Type	Creation date	Publisher
No records found.				

Save Save and close Revert Close

Tags

Applied tags

Tag A Tag H

Apply tag

Search or create tag

Last used click to apply

Tag B

Most used click to apply

Tag B Tag C

My tags click to apply

Tag B Tag C

Applying tags

Applying existing tags

Apply a tag from the sections **Last used**, **Most used** and **My tags** by clicking on it.

Master Data - Reference ▾ Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Tutorial ▾

Actions ▾

Titles

Publishers

▼ Authors

lastNameAuthor1 firstNameAuthor1

Royalties

Contracts

BookRental

Authors EditorWithContracts

Id 00-0-001

Lastname lastNameAuthor1

Firstname firstNameAuthor1

Address addressAuthor1

Birth date / /

Death date / /

City cityAuthor1

Country France

Books + Actions No records

Id	Title	Type	Creation date	Publisher
No records found.				

Save Save and close Revert Close

Tags

Applied tags

Tag A Tag H

Apply tag

Search or create tag

Last used click to apply

Tag B

Most used click to apply

Tag B Tag C

My tags click to apply

Tag B Tag C

If the required tag is not displayed under the sections **Last used**, **Most used** and **My tags** then search for it in the table's tags catalog using the input field **Search or create**.

Master Data - Reference ▾ Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Tutorial ▾

Actions ▾

Titles

Publishers

▼ Authors

lastNameAuthor1 firstNameAuthor1

Royalties

Contracts

BookRental

Authors EditorWithContracts

Id 00-0-001

Lastname * lastNameAuthor1

Firstname * firstNameAuthor1

Address addressAuthor1

Birth date / /

Death date / /

City cityAuthor1

Country France

Books + Actions ▾ No records ▾

Id	Title	Type	Creation date	Publisher
No records found.				

Save Save and close Revert Close

Tags

Applied tags

Tag A

Apply tag

Tag

Tag B

Tag C

Most used click to apply

Tag B Tag C

My tags click to apply

Tag B Tag C

For details about managing tag catalogs, see [Managing tags and tag catalogs](#).

Applying a new tag

If the required tag doesn't exist in the table's tags catalog, then create it by typing its name into the input field **Search or create** and pressing Enter or by clicking the **Send** button.

If you are an administrator for the table's tag catalog, you can choose between making the tag public or private when creating a new tag. All profiles can see public tags, and only creators can see private tags.

Master Data - Reference ▾ Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Tutorial ▾

Actions ▾

Titles

Publishers

▾ Authors

lastNameAuthor1 firstNameAuthor1

Royalties

Contracts

BookRental

Authors EditorWithContracts

Id 00-0-001

LastName * lastNameAuthor1

Firstname * firstNameAuthor1

Address addressAuthor1

Birth date / /

Death date / /

City cityAuthor1

Country France

Books + Actions ▾ No records ▾

Id	Title	Type	Creation date	Publisher
No records found.				

Save Save and close Revert Close

Tags

Applied tags

Tag A

Apply tag

Tag D

Create public tag

Create private tag

Last used

Tag B

Most used click to apply

Tag B Tag C

My tags click to apply

Tag B Tag C

For details about managing tag catalog administrators, see [Managing tags and tag catalogs](#).

Removing a tag

Remove a tag assigned to the record using the 'X' button that appears on the tag under the section **Applied tags**.

Master Data - Reference ▾ Author [lastNameAuthor1 firstNameAuthor1 [00-0-001]]

Tutorial ▾

Actions ▾

Titles

Publishers

▾ Authors

lastNameAuthor1 firstNameAuthor1

Royalties

Contracts

BookRental

Authors EditorWithContracts

Id 00-0-001

LastName * lastNameAuthor1

Firstname * firstNameAuthor1

Address addressAuthor1

Birth date / /

Death date / /

City cityAuthor1

Country France

Books + Actions ▾ No records ▾

Id	Title	Type	Creation date	Publisher
No records found.				

Save Save and close Revert Close

Tags

Applied tags

Tag A Tag H

Apply tag

Search or create tag

Last used click to apply

Tag B

Most used click to apply

Tag B Tag C

My tags click to apply

Tag B Tag C

See also:

- [Comments](#)
- [Ratings](#)

Workflow models

- [Introduction to workflow models](#)
- [Creating and implementing a workflow model](#)
- [Configuring the workflow model](#)
- [Publishing workflow models](#)

Introduction to workflow models

- [Overview](#)
 - [What is a workflow model?](#)
 - [Basic concepts related to workflow models](#)
- [Using the Workflow Models area user interface](#)
- [Generic message templates](#)
 - [Example](#)
- [Limitations of workflows](#)

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 \(glossary\)](#)
- [Data workflow \(glossary\)](#)

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](#)
- [user task](#)
- [work item](#)
- [workflow condition](#)
- [sub-workflow invocation](#)
- [wait task](#)
- [data context](#)

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.

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:

system.time	System time of the repository.
system.date	System date of the repository.
workflow.lastComment	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).
workflow.lastDecision	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).
user.fullName	Full name of the notified user.
user.login	Login of the notified user.
workflow.process.label	Label of the current workflow.
workflow.process.description	Description of the current workflow.
workflow.workItem.label	Label of the current work item.
workflow.workItem.description	Description of the current work item.

workflow.workItem.offeredTo	Role to which the current work item has been offered.
workflow.workItem.allocatedTo	User to whom the current work item has been allocated.
workflow.workItem.link	Link to access the current work item in the work item inbox, using the Web Component API. 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.
workflow.workItem.link.allocateAndStart	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. 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.
workflow.currentStep.label	Label of the current step.
workflow.currentStep.description	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

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.
- **User profile names can be up to 64 characters**, otherwise compatibility is not guaranteed.

Related concepts:

- [Data workflows](#)

Creating and implementing a workflow model

- [Creating a workflow model](#)
- [Implementing the steps](#)
 - [Progress strategy of the next step](#)
 - [Hidden in progress view](#)
- [User tasks](#)
 - [Mode](#)
 - [List of profiles](#)
 - [Service](#)
 - [Configuration](#)
- [Script tasks](#)
 - [Library script tasks](#)
 - [Specific script tasks](#)

- Conditions
 - Library conditions
 - Specific conditions
- Sub-workflow invocations
- Wait tasks
- Editing the workflow diagram
 - About
 - Actions
 - View
 - Buttons
 - Edit
 - Features

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.

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 user task with multiple work items), 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 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.

In the case of sub-workflows invocation, a dedicated property « Foreground sub-workflow » is available to precise the progress strategy. For the previous progress strategy « Open directly the next step », a foreground sub-workflow must be selected. Only the steps and associated progress strategy included in this foreground sub-workflow will be evaluated. Please note the following specific rules about the foreground sub-workflow property :

- If only one sub-workflow is launched, it is automatically considered in the foreground,
- The foreground sub-workflow property will be ignored if the previous step has not selected the progress strategy « Automatically open the next step »,
- When all the sub-workflow are completed, and if the last completed sub-workflow is the foreground one, then the progress strategy defined on the sub-workflow invocation step is evaluated: if the progress strategy is « automatically open the next step », the next step can be opened without displaying the inbox. In all cases, the progress strategy of the last step of the sub-workflow is always ignored.

Hidden in progress 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.

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](#).

List of profiles

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

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.

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](#)

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](#)

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.

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.handleWorkItemCompletion` in the JavaDoc for more information

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 `UserTask` 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](#)

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.

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:

Library script task	EBX® includes a number of built-in library script tasks, which can be used as-is. 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.
Specific script task	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.

Packaging TIBCO EBX® modules

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 `ScriptTaskBean`. 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.executeScript` is called when the data workflow reaches the corresponding step.

Attention:

The method `ScriptTaskBean.executeScript` 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](#).

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](#).

The method `ScriptTask.executeScript` is called when the data workflow reaches the corresponding step.

Attention:

The method `ScriptTask.executeScript` 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](#).

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.

Conditions

Conditions are decision steps in workflows.

Two types of conditions exist, which, once defined, can be used in workflow model steps:

Library condition	EBX® includes a number of built-in library conditions, which can be used as-is. 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.
Specific condition	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.

Packaging TIBCO EBX® modules

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 `ConditionBean`. 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](#).

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 `Condition`.

See the [example](#).

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.

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.

The property « Foreground sub-workflow » is useful to precise the progress strategy when several sub-workflows are launched. If the previous step progress strategy is « directly open the next step », this property defines which sub-workflow should be considered in the foreground (only one sub-workflow can hold this behavior). Only the work items which have been generated in this foreground sub-workflow will be able to be opened automatically without passing by the inbox. If

only one sub-workflow is defined, this property is ignored (the single sub-workflow is automatically considered in the foreground). For further information, refer to the progress strategy: [Progress strategy of the next step](#)

Two types of sub-workflow invocations exist:

Static	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. This mode is useful when the sub-workflows to be launched and the output mappings are predetermined.
Dynamic	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. The workflow data context is directly accessible from the Java bean. Dynamic sub-workflow invocations must be declared in a <code>module.xml</code> file. 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.

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: `WaitTaskBean`.

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.

Editing the workflow diagram

About

A workflow model is displayed in a BPMN-like editable diagram.

This view provides full editing capabilities and 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

Export as PNG	Creates a PNG image.

Export as SVG	Creates an SVG image.
Export as PDF	Creates a PDF document from the workflow model. Several properties are available to custom your export. You can configure the orientation of the pdf pages (Landscape or Portrait), the size (A3, Letter, etc ...). There is also two properties that might provide more details to your model: • Display identifiers (default value is true): Same behavior with "Show steps identifiers" in the workflow diagram toolbar but restricted to the pdf export. It will show the step identifiers for each step. • Add index (default value is true): At the end of the pdf, an index recording all the steps of the diagram will be added. It might be useful in the case where several steps title are too long to be fully displayed. <i>For this property, we recommend showing the steps identifiers as it will help in recognizing which index line corespond to which step.</i>

View

Layout > Default layout	Applies the default layout to the diagram.
Display > Show/Hide grid	Shows the grid if the grid is not visible, hides it otherwise.
Display > Zoom In	A zoom in is executed on the diagram. To go faster you can use "Ctrl, Up mouse wheel"

	with the mouse pointing the diagram boundary.
Display > Zoom Out	A zoom out is executed on the diagram. To go faster you can use "Ctrl, Down mouse wheel" with the mouse pointing the diagram boundary.
Display > Show/Hide steps identifiers	Show/hide a badge displaying the identifier of a workflow step.
Plan view > Hierarchy	Shows the hierarchical view of a given workflow model if enabled.

Buttons

Save layout	Saves the current layout.
Save layout and close	Saves the current layout and closes the service.
Revert	Reverts changes and reloads a previously saved layout.
Close	Closes the service.

Edit

Create	Hovering over a link or selecting it, makes '+' appear. It is used to create a new step.
Delete	DEL Pressing 'DEL' key will remove the selected node. Please note that this action cannot be undone.

Node toolbar	Hovering on a node or directly selecting it makes a toolbar appear. It is used to either: • Edit a step. • Remove a step. • Duplicate a step. • Relink to an existing step. • Show or hide in progress view.
Edit a node	Double click on a node in order to quick edit a step.

Features

The diagram view offers useful additional features

Undo last action	CTRL + Z Please note that it cannot be done after the following actions: • Create/Remove/Edit/Duplicate a step • Relink • Show/Hide in progress view • Revert
Zoom in/Zoom out.	Mouse middle button then mouse wheel / CTRL then mouse wheel. Also see View
Multiple selection	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).
Customizing links drawing	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.
Edit a step	Double clicking on a step will display an edition form.
Overview	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 for the workflow diagram view.

Configuring the workflow model

- [Information associated with a workflow model](#)
- [Workflow model properties](#)
- [Data context](#)
- [Custom workflow execution views](#)
- [Permissions on associated data workflows](#)
- [Workflow model snapshots](#)
- [Deleting a workflow model](#)

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'](#) menu for your workflow model in the navigation pane.

Owner	Specifies the workflow model owner, who will have the rights to edit the workflow model's information and define its permissions.
Localized documentation	Localized labels and descriptions for the workflow model.
Activated	<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.

Workflow model properties

Configuration for a workflow model is accessible in the navigation pane under 'Workflow model configuration'.

Module name	Module containing specific Java resources (user task extensions, specific scripts and conditions).
Notification of start	The list of profiles to which to send notifications, based on a template, when a data workflow is launched. See Generic message templates.
Notification of completion	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. See Generic message templates.
Notification of error	The list of profiles that will receive notifications, based on a template, when a data workflow is in error state. See Generic message templates.
Priority	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 Work item priorities for more information. Note: Only users who are defined as workflow administrators will be able to manually modify the priority level of any associated data workflows.
Activate quick launch	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.
Automatically open the first step	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.
Workflow trigger	Component that intercepts the main events of a workflow. This bean must be declared in a <code>module.xml</code> file. See the example.
Permissions	Permissions on actions related to the data workflows associated with the workflow model. This bean must be declared in a <code>module.xml</code> file. See the example.
Programmatic action permissions	Defines a custom component that handles the permissions of the workflow. If set, this overrides all permissions defined in the property 'Permissions'.

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:

Name	Identifier of the variable.
Default value	If defined, the variable will be initialized with this default value.
Input parameter	'Yes' must be checked in order to define this variable as an input parameter.

Output parameter	'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.

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.

Permissions on associated data workflows

Workflow administration	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.
Workflow administration > Replay a step	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.

Workflow administration > Terminate workflow	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.
Workflow administration > Force a workflow to resume	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.
Workflow administration > Disable a publication	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.
Launch workflows	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. Permissions to launch a workflow are always automatically verified, even by API. It is possible to bypass these automatic

	permission checks and manage specific authorization rules manually. For this, set the <code>ebx.properties</code> <code>permissionCheckDelegated</code> to true and use the dedicated API <pre>WorkflowEngine.isUserAuthorizedToLaunch</pre> to delegate permission evaluation to the developer.
Visualize workflows	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.
Visualize workflows > The workflow creator can visualize it	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'.
Visualize workflows > Visualize completed workflows	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](#)

Workflow model snapshots

The history of workflow model snapshots can be managed from **Actions > View publications 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.

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](#)

Publishing workflow models

- [About workflow publications](#)

- [Publishing and workflow model snapshots](#)
- [Sub-workflows in publications](#)

About workflow publications

Once a workflow model is defined, it must be published in order to enable authorized users to launch associated data workflows.

Two services are available in order to publish :

- With the button **Publish** in the navigation pane: this service allows to publish only the model currently edited.
- Through the button **Mass publish** that can be found in the toolbar on the same level of the section title **Workflow models** : this service allows many models to be published at once. The first step allows to select many models to publish: only models not containing sub-workflows invocation will appear in this step. Indeed, a workflow having sub-workflows doesn't qualify for mass publish.

Workflow models can be published several times. A publication is identified by its publication name

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](#)

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

- [Introduction to data workflows](#)
- [Using the Data Workflows area user interface](#)
- [Work items](#)
- [Managing data workflows](#)
 - [Launching and monitoring data workflows](#)
 - [Administration of data workflows](#)

Introduction to data workflows

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](#)
- [Launching and monitoring data workflows](#)
- [Administration of data workflows](#)
- [Permissions on associated data workflows](#)

Related concepts:

- [Workflow models](#)


Using the Data Workflows area user interface

- [Navigating within the interface](#)
- [Navigation rules](#)
 - [Work items inbox](#)
 - [Workflow launchers](#)
- [Custom views](#)

- Specific columns
- Filtering items in views
- Workflow progress view

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:

Work items inbox	All work items either allocated or offered to you, for which you must perform the defined task.
Workflow launchers	List of workflow model publications from which you are allowed to launch data workflows, according to your user permissions.
Monitoring	Monitoring views on the data workflows for which you have the necessary viewing permissions.
Publications	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.
Active workflows	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.
Work items	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.
Completed workflows	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.

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](#) 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](#) in workflow modeling.

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](#) for more information.


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.

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.



Workflow progress 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 progress 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).

Work items

- [About work items](#)
 - [Work item states](#)
- [Working on work items as a participant](#)
- [Work item priorities](#)

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](#)

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





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 progress 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.

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 \(glossary\)](#)

Related concepts:

- [User tasks](#)

Managing data workflows

- [Launching and monitoring data workflows](#)
- [Administration of data workflows](#)

Launching and monitoring data workflows

- [Launching data workflows](#)
- [Monitoring activities](#)
- [Managing work item allocation](#)

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.

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.

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.

Allocate	Allocate a work item to a specific user. This action is available for work items in the <i>offered</i> state.
Deallocate	Reset a work item in the <i>allocated</i> state to the <i>offered</i> state.

Reallocate

Modify the user to whom a work item is allocated. This action is available for work items in the *allocated* state.

See also:

- [Work items](#)
- [Permissions on associated data workflows](#)

Related concepts:

- [Workflow models](#)

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.

- [Overview of data workflow execution](#)
- [Data workflow administration actions](#)
 - [Actions on publications](#)
 - [Actions on data workflows](#)

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](#)

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](#).
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](#)

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](#)

Data services

- [Introduction to data services](#)
- [Generating data service WSDLs](#)

Introduction to data services

- [Overview](#)
 - [What is a data service?](#)
 - [Lineage](#)
 - [Glossary](#)
- [Using the Data Services area user interface](#)

Overview

What is a data service?

A [data service](#) 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 RESTful service that allows interaction with EBX®.

The built-in REST data services have an OpenAPI interface, they are also 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](#)
- [REST](#)

Lineage


[Lineage](#) 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](#)

Using the Data Services area user interface



Note:

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

Related concepts:

- [Dataspace](#)
- [Dataset](#)
- [Data workflows](#)
- [Introduction](#)

Generating data service WSDLs

- [Generating a WSDL for operations on data](#)

- Operations on datasets
- Operations on tables
- Generating a WSDL for dataspace operations
 - Operations on dataspace
- Generating a WSDL for data workflow operations
 - Operations on data workflows
- Generating a WSDL for lineage
- Generating a WSDL for administration
 - Operations for administration
- Generating a WSDL to modify the default directory
 - Operations on the default directory

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 dataspace 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](#)
- [Operations generated from a data model](#)

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](#)
- [Operations on datasets and dataspace](#)

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](#)
- [Operations on data workflows](#)

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](#).

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](#)

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](#)
- [User interface operations](#)
- [System information operation](#)

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](#).

See also:

- [WSDL download from HTTP protocol](#)
- [Directory services](#)

Reference Manual

- Integration
 - Overview of integration and extension
 - Using TIBCO EBX® as a Web Component
 - Built-in user services
 - Supported XPath syntax
- Localization
 - Labeling and localization
 - Extending TIBCO EBX® internationalization
- Persistence
 - Overview of persistence
 - History
 - Replication
 - Data model evolutions
- Other
 - Inheritance and value resolution
 - Permissions
 - Criteria editor
 - Search
 - Performance and tuning

Integration

- Overview of integration and extension
- Using TIBCO EBX® as a Web Component
- Built-in user services
- Supported XPath syntax

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.

- [User interface customization and integration](#)
- [Data services](#)
- [XML and CSV import/export services](#)
- [Programmatic services](#)

User interface customization and integration

The EBX® graphical interface can be customized through various EBX® APIs.

It can also be integrated into any application that is accessible through a [Supported web browsers](#).

See [Interface customization](#) for more information.

Data services

The data services 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](#)
- [REST data services](#)

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
- the data services
- the Java API

See also:

- [XML import and export](#)
- [CSV import and export](#)

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](#) are not sufficient.

See also:

- `ProgrammaticService`

Using TIBCO EBX® as a Web Component

- [Overview](#)
- [Integrating EBX® Web Components into applications](#)
- [Repository element and scope selection](#)
- [Combined selection](#)
 - [Specific case](#)
 - [Known limitations](#)
- [Request specifications](#)

- Base URL
 - User authentication and session information parameters
 - Entity and service selection parameters
 - Layout parameters
- [Example calls to an EBX® Web Component](#)

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](#)

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](#).

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](#) or a [Custom widget](#), 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 `UIHttpManagerComponent` in the API.

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](#)

See also:

- [firstCallDisplay](#)

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](#)

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.

Request specifications

Base URL

In a default deployment, the base URL must be of the following form:

```
http://[:]/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
<code>login</code> and <code>password</code> , 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. See <code>Directory</code> for more information.	No
<code>trackingInfo</code>	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. See <code>AccessRule</code> for more information.	No
<code>redirect</code>	The URL to which the user will be redirected at the end of the component session,	No

PARAMETER	DESCRIPTION	REQUIRED
	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> . For more information, see Exit policy .	
<code>locale</code>	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
<code>branch</code>	Selects the specified dataspace.	No
<code>version</code>	Selects the specified snapshot.	No
<code>instance</code>	Selects the specified dataset. The value must be the reference of a dataset that exists in the selected dataspace or snapshot.	Only if <code>xpath</code> or <code>viewPublication</code> is specified.
<code>viewPublication</code>	Specifies the publication name of the tabular, hierarchical, or tile view to apply to the selected content. 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</i> > <i>Views</i> . All settings of the	No

PARAMETER	DESCRIPTION	REQUIRED
	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 <code>xpath</code> parameter as a logical 'AND' operation.	
<code>xpath</code>	Specifies a node selection in the dataset. Value may be a valid absolute path located in the selected dataset. The notation must conform to a simplified XPath, with abbreviated syntax. It can also be a predicate surrounded by "" (to be encoded using %5B and %5D respectively) if a table can be automatically selected using other Web Component parameters (for example, <code>viewPublication</code> or <code>workflowView</code>). For XPath syntax, see XPath supported syntax . See <code>UIHttpManagerComponent.setPredicate</code> for more information.	No
<code>service</code>	Specifies the service to access. For more information on built-in User services, see Built-in services . In the Java	No

PARAMETER	DESCRIPTION	REQUIRED
	API, see <code>ServiceKey</code> for more information.	
<code>workflowView</code>	Specifies the workflow section to be selected. See <code>WorkflowView</code> for more information.	No.
<code>perspectiveName</code>	Specifies the name of the perspective to be selected. If this parameter is specified, the <code>scope</code> parameter can have only two values: full and data.	Only if <code>perspectiveActionId</code> or <code>perspectiveActionName</code> is specified.
<code>perspectiveActionId</code>	Deprecated. Please consider using <code>perspectiveActionName</code> instead. Specifies the identifier of the perspective action to be selected.	No.
<code>perspectiveActionName</code>	Specifies the unique name of the perspective action to be selected.	No.

Layout parameters

PARAMETER	DESCRIPTION	REQUIRED
<code>scope</code>	Specifies the scope to be used by the web component. Value can be <code>full</code> , <code>data</code> , <code>dataspace</code> , <code>dataset</code> or <code>node</code> .	No, default will be computed to be the smallest possible according to the target selection.

PARAMETER	DESCRIPTION	REQUIRED
	See <code>UIHttpManagerComponent.Scope</code> for more information.	
<code>firstCallDisplay</code>	Specifies which display must be used instead of the one determined by the combination of selection and <code>scope</code> parameter. Possible values are: • <code>auto</code>: The display is automatically set according to the selection. • <code>view</code>: Forces the display of the tabular, the hierarchical or the tile view. • <code>record</code>: If the predicate has at least one record, forces the display of the first record in the list. For example, <code>firstCallDisplay=view</code> <code>firstCallDisplay=view:hierarchyExpanded</code> <code>firstCallDisplay=record</code> <code>firstCallDisplay=record:{predicate}</code>	No, default will be computed according to the target selection.

PARAMETER	DESCRIPTION	REQUIRED
	See <code>UIHttpManagerComponent.setFirstCallDisplay</code> for more information. See <code>UIHttpManagerComponent.setFirstCallDisplayHierarchyExpanded</code> for more information. See <code>UIHttpManagerComponent.setFirstCallDisplayRecord</code> for more information.	
<code>closeButton</code>	Specifies how to display the session close button. Value can be <code>logout</code> or <code>cross</code>. See <code>UIHttpManagerComponent.CloseButtonSpec</code> for more information.	No. If scope is not <code>full</code> , no close button will be displayed by default.
<code>dataSetFeatures</code>	Specifies which features to display in a UI service at the dataset level or a form outside of a table. These options pertain only to features in the workspace. It is recommended to use this property with the smallest <code>scope</code> possible, namely <code>dataset</code> or <code>node</code>. Syntax:	No.

PARAMETER	DESCRIPTION	REQUIRED
	<code><prefix> ":" <feature></code> <code>[", " <feature>]*</code> where <code><prefix></code> is <code>hide</code> or <code>show</code>, <code><feature></code> is <code>services</code>, <code>title</code>, <code>save</code>, or <code>revert</code>. For example, <code>hide:title</code> <code>show:save, revert</code> See <code>UIHttpManagerComponent.DataSetFeatures</code> for more information.	
<code>viewFeatures</code>	Specifies which features to display in a tabular, a hierarchy or a tile view (at the table level). These options pertain only to features in the workspace. It is recommended to use this property with the smallest <code>scope</code> possible, namely <code>dataset</code> or <code>node</code>. Syntax: <code><prefix> ":" <feature></code> <code>[", " <feature>]*</code>	No.

PARAMETER	DESCRIPTION	REQUIRED
	where • <code><prefix></code> is <code>hide</code> or <code>show</code>, • <code><feature></code> is <code>create</code>, <code>views</code>, <code>selection</code>, <code>filters</code>, <code>services</code>, <code>refresh</code>, <code>title</code>, or <code>breadcrumb</code>. For example, <code>hide:title,selection</code> <code>show:service,title,breadcrumb</code> See <code>UIHttpManagerComponent.ViewFeatures</code> for more information.	
<code>recordFeatures</code>	Specifies which features must be displayed in a form at the record level. These options pertain only to features in the workspace. It is recommended to use this property with the smallest <code>scope</code> possible, namely <code>dataset</code> or <code>node</code>. Syntax: <code><prefix> ":" <feature></code> <code>["," <feature>]*</code> where	No.

PARAMETER	DESCRIPTION	REQUIRED
	• <code><prefix></code> is <code>hide</code> or <code>show</code> , • <code><feature></code> is <code>services</code> , <code>title</code> , <code>breadcrumb</code> , <code>save</code> , <code>saveAndClose</code> , <code>close</code> , or <code>revert</code> . For example, <code>hide:title</code> <code>show:save,saveAndClose,revert</code> See <code>UIHttpManagerComponent.RecordFeatures</code> for more information.	
<code>pageSize</code>	Specifies the number of records that will be displayed per page in a table view (either tabular, hierarchical or tile).	No.
<code>startWorkItem</code>	Specifies a work item must be automatically taken and started. Value can be <code>true</code> or <code>false</code> . See <code>ServiceKey.WORKFLOW</code> for more information.	No. Default value is <code>false</code> , where the target work item state remains unchanged.

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 'instanceld' 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"]
```

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
- when defining perspective action menu items
- as extended user services when used with service extensions
- when using EBX® as a Web Component

This reference page describes the built-in user services and their parameters.

- Access data (default service)
 - Input parameters
- Create a new record
 - Input parameters
 - Output parameters
- Duplicate a record
 - Input parameters
 - Output parameters
- Export data to an XML file
 - Input parameters
- Export data to a CSV file
 - Input parameters
- Import data from an XML file
 - Input parameters
- Import data from a CSV file

- Input parameters
- Access a dataspace
 - Input parameters
- Validate a dataspace, a snapshot or a dataset
 - Input parameters
 - Output parameters
- Merge a dataspace
 - Input parameters
 - Output parameters
- Access the dataspace merge view
 - Input parameters
- Compare contents
 - Input parameters
- Data workflows
 - Input parameters

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace is required for this service.
<code>disableAutoComplete</code>	Disable Accept at start	By default, the interaction associated with this service is

PARAMETER	LABEL	DESCRIPTION
		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.
<code>firstCallDisplay</code>	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.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace.
<code>scope</code>	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

PARAMETER	LABEL	DESCRIPTION
		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.
<code>trackingInfo</code>	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.
<code>version</code>	Snapshot	The identifier of the specified snapshot - A dataspace is required for this service.
<code>viewPublication</code>	View	The publication name of the view to display. The view must be configured for the selected table.
<code>xpath</code>	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.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - This field is required for this service.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.
<code>scope</code>	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.
<code>trackingInfo</code>	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.
<code>xpath</code>	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
<code>created</code>	Created record	Contains the XPath of the created record.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - This field is required for this service.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.
<code>scope</code>	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.
<code>trackingInfo</code>	Tracking information	Tracking information is logged into 'history' logs. It

PARAMETER	LABEL	DESCRIPTION
		may also be used for any other purpose like access control or additional export information.
<code>xpath</code>	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
<code>created</code>	Created record	Contains the XPath of the created record.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace is required for this service.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.

PARAMETER	LABEL	DESCRIPTION
<code>scope</code>	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.
<code>trackingInfo</code>	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.
<code>version</code>	Snapshot	The identifier of the specified snapshot - A dataspace is required for this service.
<code>xpath</code>	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.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace is required for this service.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.
<code>scope</code>	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.
<code>trackingInfo</code>	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.
<code>version</code>	Snapshot	The identifier of the specified snapshot - A dataspace is required for this service.
<code>xpath</code>	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

PARAMETER	LABEL	DESCRIPTION
		abbreviated syntax - This field is required for this service.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace is required for this service.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.
<code>scope</code>	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.
<code>trackingInfo</code>	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.

PARAMETER	LABEL	DESCRIPTION
<code>xpath</code>	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.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace is required for this service.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace - This field is required for this service.
<code>scope</code>	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.

PARAMETER	LABEL	DESCRIPTION
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.

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.

PARAMETER	LABEL	DESCRIPTION
<code>trackingInfo</code>	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.
<code>version</code>	Snapshot	The identifier of the specified snapshot - A dataspace is required for this service.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace or snapshot is required for this service.
<code>instance</code>	Dataset	The value must be the reference of a dataset that exists in the selected dataspace.
<code>scope</code>	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

PARAMETER	LABEL	DESCRIPTION
		depends on the selected entities or service.
<code>trackingInfo</code>	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.
<code>version</code>	Snapshot	The identifier of the specified snapshot - A dataspace or snapshot is required for this service.

Output parameters

PARAMETER	LABEL	DESCRIPTION
<code>hasError</code>	Found errors	Contains 'true' if validation has produced errors.
<code>hasFatal</code>	Found fatal errors	Contains 'true' if validation has produced fatal errors.
<code>hasInfo</code>	Found informations	Contains 'true' if validation has produced informations.
<code>hasWarning</code>	Found warnings	Contains 'true' if validation has produced warnings.

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
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace is required for this service.
<code>scope</code>	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.
<code>trackingInfo</code>	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
<code>mergeResult</code>	Merge success	Contains 'true' if merge succeeded, otherwise 'false'.
<code>mergeState</code>	Merge state	Contains the return code of the merge. It is strongly recommended to parse this value by using the <code>InteractionMergeState</code> <code>UIHttpManagerComponentReturnCode</code> .

Access the dataspace merge view

The merge.view service is automatically considered complete.

Service name parameter: `service=@merge.view`

Input parameters

PARAMETER	LABEL	DESCRIPTION
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace is required for this service.
<code>scope</code>	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.
<code>trackingInfo</code>	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.

Compare contents

Workflows automatically consider the compare service as complete.

Service name parameter: `service=@compare`

Input parameters

PARAMETER	LABEL	DESCRIPTION
<code>branch</code>	Dataspace	The identifier of the specified dataspace - A dataspace or snapshot and a dataspace or snapshot to compare to are required for this service.
<code>compare.branch</code>	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.
<code>compare.filter</code>	Comparison filter	To ignore inheritance and function 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.
<code>compare.instance</code>	Dataset to compare	The value must be the reference of a dataset that exists in the selected dataspace to compare.
<code>compare.version</code>	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.
<code>compare.xpath</code>	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

PARAMETER	LABEL	DESCRIPTION
		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 conform to a simplified XPath, in its abbreviated syntax.

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
<code>scope</code>	Scope	Defines the scope of the user navigation for this service.
<code>viewPublication</code>	View publication	Defines the publication name of the view to apply for this service.
<code>workflowView</code>	Workflow view	Specifies the workflow view type. Value can be one of the following: "inbox", "launcher", "monitoringPublications", "monitoringWorkflows", "monitoringWorkItems" or "completedWorkflows".
<code>xpath</code>	Filter (XPath)	An optional filter. The syntax should conform to an XPath predicate surrounded by "".

Supported XPath syntax

- [Overview](#)
- [Example expressions](#)
 - [Absolute path](#)
 - [Relative paths](#)
 - [Root and descendant paths](#)

- Table paths with predicates
- Complex predicates
- Predicates with parameters
- Search predicate
- Predicates for validation search
- Syntax specifications for XPath expressions
 - Overview
 - Predicate specification
 - Extraction of foreign keys
- Java API

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.

Example expressions

The general XPath expression is:

```
path
```

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.setXPathParameter` of the Java API.

Search predicate

- Syntax: `osd:search(fields, queryString[, templateKey])`
- Examples:
 - `osd:search('col1', 'xxx')`
 - `osd:search('col1,col2', 'xxx')`
 - `osd:search('', 'xxx')`
 - `osd:search(col1, 'xxx', myTemplate@myModule)`

The `osd:search` function tries to match a term, or a list of terms, against the set of fields of the current table. This function is generic, handling every field datatype supported by EBX®. When no fields are specified, it searches against all fields for current table. For a more advanced usage, the query string supports specialized operators, see [Special characters](#) for more information.

For any concerned field, if a label exists, the search targets the label, rather than the value; however, this is not yet supported in some cases. See [Limitations](#) for more information.

The predicate `osd:search` is localized.

Note:

The locale can be set by the methods of the Java API `Request.setLocale` or `Request.setSession`.

Note:

The identifier of a [search template](#) can be specified, to customize the behavior of the search.

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')
```

- 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.

Note:

The locale can be set by the methods of the Java API `Request.setLocale` or `Request.setSession`.

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.

Syntax specifications for XPath expressions

Overview

EXPRESSION	FORMAT	EXAMPLE
XPath expression	<i><container path></i>	<code>/books[title='xxx']</code>

EXPRESSION	FORMAT	EXAMPLE
<container path>	<absolute path> or <relative path>	
<absolute path>	/a/b or //b	//books
<relative path>	../..b , ./b or b	../..books

Predicate specification

EXPRESSION	FORMAT	NOTES/EXAMPLE
<predicate>	Example: A and (B or not(C)) A,B,C: <atomic expression>	Composition of: logical operators parentheses, not() and atomic expressions.
<atomic expression>	<path><comparator><criterion> or method(<path>,<criterion>)	royalty = 24.5 starts-with(title, 'Johnat') booleanValue = true
<path>	<relative path> or osd:label(<relative path>)	Relative to the table that contains it: ../authorstitle
<comparator>	<boolean comparator> <numeric comparator> or <string comparator>	
<boolean comparator>	= or !=	
<numeric comparator>	= , != , < , > , <= , or >=	

EXPRESSION	FORMAT	NOTES/EXAMPLE
<i><string comparator></i>	= or !=	
<i><method></i>	• <i><date method></i> • <i><string method></i> • <code>osd:is-null</code> method or • <code>osd:is-not-null</code> method	
<i><date, time & dateTime method></i>	• <code>date-less-than</code> • <code>date-equal</code> or • <code>date-greater-than</code>	
<i><string method></i>	• <code>contains</code> • <code>osd:contains-case-insensitive</code> • <code>osd:contains-record-label</code> • <code>osd:is-equal-case-insensitive</code> • <code>osd:is-empty</code> • <code>osd:is-empty-or-nil</code> • <code>osd:is-neither-empty-nor-nil</code> • <code>osd:is-not-empty</code> • <code>ends-with</code> • <code>osd:ends-with-case-insensitive</code>	

EXPRESSION	FORMAT	NOTES/EXAMPLE
	• <code>matches</code> • <code>osd:search</code> • <code>starts-with</code> or • <code>osd:starts-with-case-insensitive</code>	
<code><criteriaion></code>	• <code><boolean criteriaion></code> • <code><date criteriaion></code> • <code><dateTime criteriaion></code> • <code><numeric criteriaion></code> • <code><string criteriaion></code> or • <code><time criteriaion></code>	
<code><boolean criteriaion></code>	<code>true</code> , <code>false</code>	
<code><numeric criteriaion></code>	An integer or a decimal	<code>-4.6</code>
<code><string criteriaion></code>	Quoted character string	<code>'azerty'</code>
<code><date criteriaion></code>	Quoted and formatted as 'yyyy-MM-dd'	<code>'2007-12-31'</code>
<code><time criteriaion></code>	Quoted and formatted as 'HH:mm:ss' or 'HH:mm:ss.SSS'	<code>'11:55:00'</code>
<code><dateTime criteriaion></code>	Quoted and formatted as 'yyyy-MM-ddTHH:mm:ss' or 'yyyy-MM-ddTHH:mm:ss.SSS'	<code>'2007-12-31T11:55:00'</code>

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`

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. The criterion value must be a single value. That means that aggregated lists are not allowed as the target of the contextual path.

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.

In addition, the predicates `osd:is-empty` and `osd:is-not-empty` can be applied on any node which type is a list to verify if the list is empty or not.

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.encodeLiteralStringWithDelimiters` in the Java API handles this.

Examples of using `encodeLiteralStringWithDelimiters`

VALUE OF LITERAL EXPRESSION	RESULT OF THIS METHOD
<code>Coeur</code>	<code>'Coeur'</code>
<code>Coeur d'Alene</code>	<code>"Coeur d'Alene"</code>
<code>He said: "They live in Coeur d'Alene".</code>	<code>'He said: "They live in Coeur d'Alene".'</code>

Extraction of foreign keys

In EBX®, the foreign keys are grouped into a single field with the `osd:tableRef` 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](#) 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).

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

- [Labeling and localization](#)
- [Extending TIBCO EBX® internationalization](#)

Labeling and localization

- [Overview](#)
 - [Localizing user interactions](#)
 - [Textual information](#)
- [Value formatting policies](#)
- [Syntax for locales](#)

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. Data model: If a data model has been localized into other languages than those natively supported by EBX®, the user can select one of those languages for that particular data model. See [Extending TIBCO EBX® internationalization](#) 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 by the module.

See also:

- [Labels and messages](#)
- [Tables declaration](#)
- [Foreign keys declaration](#)

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](#).

For each locale declared by the module, 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 in the module, 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`.

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](#)

Extending TIBCO EBX® internationalization

- [Overview of the native EBX® localization](#)
- [Language packs software](#)
- [Extending EBX® user interface localization](#)
 - [Adding a new locale](#)
- [Localized resources resolution](#)
- [Known limitations](#)
 - [Non extendable materials](#)

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.

Language packs software

Language packs for TIBCO EBX® and its addons are available from [TIBCO eDelivery](#) according to the EBX® release. A language pack can also contain the online help, please see the `TIB_ebx-lp_<version>_readme.txt` file for more information.

Note:

On each release, there is a certain delay between the GA and the availability of the corresponding pack; in this case, the previous pack can be deployed (but may not cover all messages).

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](#).

- Package the following files in a jar file (sample name: `ebx-lp_<version>_languagepack-xx.jar`):

- 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".

A complete template set of files is available [here](#).

- Deploy the jar file in the EBX® class-path it must be accessible from the `ebx.jar` class-loader.

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.

Known limitations

Non extendable materials

Localization of the following cannot be extended:

- EBX® HTML editor and viewer.

Persistence

- [Overview of persistence](#)
- [History](#)
- [Replication](#)
- [Data model evolutions](#)

Overview of persistence

This chapter is an introduction to history tables and replicated tables.

Note:

The term *mapped mode* refers to any tables that are stored as-is, and thus whose contents can be accessed directly in the database.

- [Primary persistence of managed master data](#)
- [Historization](#)
- [Replication](#)
- [Mapped mode](#)
 - [Overview of mapped mode](#)
 - [Structural constraints](#)
 - [Data model restrictions due to mapped mode](#)

Primary persistence of managed master data

Data that is modeled in and governed by the EBX® repository are primarily persisted in the relational database, using generic tables (common to all datasets and data models).

Historization

Master data tables can activate historization in order to track modifications to their data, regardless of whether they are replicated.

The history itself is in mapped mode, meaning that it can be consulted directly in the underlying database.

See also:

- [History](#)

Replication

Replication enables direct SQL access to tables of master data, by making a copy of data in the repository to replica tables in the relational database. Replication can be enabled on any table regardless of 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](#)

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®. History tables 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](#)
- [Data model evolutions](#)

Structural constraints

When a mapped mode is set, some EBX® data model constraints will generate a "structural constraint" on the underlying RDBMS schema. This concerns the following constraining 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: when a transaction does not comply with a blocking constraint, it is cancelled and a `ConstraintViolationException` is thrown.

See also:

- [Blocking and non-blocking constraints](#)

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](#)
- 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`.
- 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, 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](#)

History

- [Overview](#)
- [Configuring history](#)
 - [Configuring history in the repository](#)
 - [Configuring history in the data model](#)
 - [Integrity](#)
- [History views and permissions](#)
 - [Table history view](#)
 - [Permissions for table history](#)
 - [Transaction history views](#)
- [SQL access to history](#)
 - [Access restrictions](#)
 - [Relational schema overview](#)
 - [Example of a generated history table](#)
 - [Combination of operations](#)
 - [Values for operation fields](#)

- Impacts and limitations of historized mode
 - Validation
 - Data model restrictions for historized tables
 - Other limitations of historized mode

Overview

History is a feature allowing to track all data modifications on a table (records creation, update and deletion).

It is an improvement over the deprecated [XML audit trail](#).

See also:

- [History](#)
- [Replication](#)
- [Data model evolutions](#)

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
<code>ebx-referenceBranch</code>	This profile is activated only on the reference dataspace.
<code>ebx-allBranches</code>	This profile is activated on all dataspaces.
<code>ebx-instanceHeaders</code>	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>

  <primaryKey>/key</primaryKey>

  <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](#) 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](#)).

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.

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](#) section. To access the table history view from Java, the method `AdaptationTable.getHistory` 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](#). The methods `Adaptation.isHistory` and `AdaptationTable.getHistory` can then be used in the programmatic rule in order to implement specific behavior for history.

Note:

There is currently a limitation when a table has a scripted permission rule on record specified: for security reason access to the table history is totally disabled for everyone. Access to history will be allowed in a future version.

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 dataspace and using the **Actions** menu in the workspace.

For more information, see [transaction history view](#).

SQL access to history

This section describes how to directly access the history data by means of SQL.

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](#).

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):

Common and generic tables	The main table is <code>HV_TX</code> ; each record of this table represents a transaction. Only transactions that involve at least one historized table are recorded. These common tables are all prefixed by "HV".
Specific generated tables	For each historized table, a specific history table is generated. This table contains the history of the data modifications on the table. 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".


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 -> C`
- `D + U -> D`
- `D + C -> U`
- `C + D -> {}` (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](#) 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.

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](#).

Data model restrictions for historized tables

Some restrictions apply to data models containing historized tables:

- [Data model restrictions due to mapped mode](#)
- 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.
- Linked fields 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](#)

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.

Replication

- [Overview](#)
- [Configuring replication](#)
 - [Enabling replication](#)
 - [Disabling replication on a specific field or group](#)
- [Accessing a replica table using SQL](#)
 - [Finding the replica table in the database](#)
 - [Access restrictions](#)
 - [SQL reads](#)
- [Requesting an 'onDemand' replication refresh](#)
- [Impact and limitations of replication](#)
 - [Validation](#)
 - [Data model restrictions for replicated tables](#)
 - [Database configuration](#)
 - [Distributed data delivery \(D3\)](#)

- [Other limitations of replication](#)

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, 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:

- [History](#)
- [Data model evolutions](#)
- [Repository administration](#)

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

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.	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: <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	Yes

ELEMENT	DESCRIPTION	REQUIRED
	corresponding insert, update, and delete statements on the replica table. <code>onDemand</code>: The replication of specified tables is only done when an explicit refresh operation is performed. See Requesting an 'onDemand' replication refresh.	
<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	Yes

ELEMENT	DESCRIPTION	REQUIRED
	be unique amongst all replications units.	

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>

        <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](#)
- 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](#)).

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.

Accessing a replica table using SQL

This section describes how to directly access a replica table using SQL.

See also:

- [SQL access to history](#)

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.

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](#)

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.

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](#) for data services for more information.
- **Java API:** Call the `ReplicationUnit.performRefresh` methods in the `ReplicationUnit` API to launch a refresh of the replication unit.

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](#) 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](#).

Data model restrictions for replicated tables

Some restrictions apply to data models containing tables that are replicated:

- [Data model restrictions due to mapped mode](#)
- 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](#) 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](#) for more information.
- Computed values are ignored.
- Linked fields 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](#)

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.performRefresh` method.

See also:

- [D3 overview](#)

Other limitations of replication

- [Limitations of supported databases](#)
- 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.

Data model evolutions

This chapter describes the modifications that are possible on data models, as well as potential limitations.

Attention:

Whenever the data modeler performs an evolution on the primary key of a table, the resulting definition is considered as a new table. 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; 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](#)
- [Types of permitted evolutions](#)
 - [Model-level evolutions](#)
 - [Table-level evolutions](#)
 - [Field-level evolutions](#)
- [Limitations/restrictions](#)
 - [Limitations related to primary key evolutions](#)

- Limitations related to foreign key evolutions
- Limitations related to field-level evolutions

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:

- 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 contains historized tables, this change marks the associated mapped tables as disabled. See [Database mapping](#) 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, this change marks the mapped table as disabled. See [Database mapping](#) for the actual removal of associated database objects.
- 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. 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 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](#) and [Disabling replication on a specific field or group](#).
- The facets of a field can be modified, except for the facets listed under [Limitations/restrictions](#).

The following 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.

Limitations/restrictions

Limitations related to primary key evolutions

When a primary key definition is modified:

- The content of the table will be reset to an empty content, in all datasets and dataspace.
- If the new primary key has been used in the past, the content of the table will be reset to the previous data existing at the time this primary key was used, in all datasets and dataspace.
- Type conversion is not supported for a primary key field. Thus, the content of the table will always be reset to empty when changing the type of a primary key even when reusing a type already used in previous primary key definition.
- The modification will be rejected if the table has - or has had - history activated in the existing dataspace. A possible workaround: first drop the history table associated with the dedicated table, then proceed to modifying the primary key. For the procedure to purge mapped table database resources, see [Database mapping](#).

Note:

If the modified primary key is referenced in the primary key of another table, all the limitations mentioned above apply to the target table.

Limitations related to foreign key evolutions

- When the declaration of a `osd:tableRef` facet is added or modified, or when the primary key of its target table is modified, the existing values will restart from empty (except if this modification is reverting to a previous definition; in this case, the previous content will be retrieved).
- In replication 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

When changing the type of a field to an incompatible type or cardinality, the field will be considered as a new one, and start with an empty content. The previous content will be retrieved if the model is rolled back to a previous definition.

- The following types are fully inter-convertible (meaning these types have the same exact persistent representation, and can be substituted to each other in the following charts):
 - `xs:string`
 - `osd:color`
 - `osd:datasetName`
 - `osd:dataspaceKey`
 - `osd:email`

- `osd:html`
 - `osd:local`
 - `osd:resource`
 - `xs:nmtoken`
 - `xs:nmtokens`
 - `osd:text`
 - `xs:anyUri`
 - `xs:name`
- The following conversions are fully supported (that is, regardless of their cardinalities):
 - `xs:decimal` to `xs:string`
 - `xs:datetime` to `xs:string`
 - `xs:date` to `xs:string`
 - `xs:integer` to `xs:string`
 - `xs:int` to `xs:decimal`
 - `xs:integer` to `xs:decimal`
 - `xs:decimal` to `xs:integer` (losing the decimal part)
 - `xs:int` to `xs:integer`
 - `xs:datetime` to `xs:date` (losing the time part)
 - `xs:date` to `xs:datetime` (defaulting the time part to 0)
 - The following conversions are possible only if the original type is single-valued:
 - `xs:boolean` to `xs:string`
 - `xs:time` to `xs:string`
 - `xs:int` to `xs:string`
 - `xs:long` to `xs:string`

The cardinality of a type can be changed; when the conversion is supported, it has the following behavior:

- 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.

Attention:

Groups and complex types do not support conversion to (and from) any other types. Moreover, when a group or complex type changes between single-occurrence and multi-occurrence, the conversion is supported only if the group or complex type is terminal.

Other

- [Inheritance and value resolution](#)
- [Permissions](#)
- [Criteria editor](#)
- [Search](#)
- [Performance and tuning](#)

Inheritance and value resolution

- [Overview](#)
 - [Dataset inheritance](#)
 - [Inherited fields](#)
 - [Computed values \(functions\)](#)
- [Dataset inheritance](#)
 - [Dataset inheritance declaration](#)
 - [Value lookup mechanism](#)
 - [Record lookup mechanism](#)
 - [Defining inheritance behavior at the table level](#)
- [Inherited fields](#)
 - [Field inheritance declaration](#)
 - [Value lookup mechanism](#)

- [Optimize & Refactor service](#)
 - [Procedure details](#)
 - [Service availability](#)

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 \(glossary\)](#)

Dataset inheritance

Dataset inheritance is particularly useful when data applies to global enterprise contexts, such as subsidiaries or business partners.

Given 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](#).

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 the associated 'FamilyOfProducts'.

Note:

When using both inheritance mechanisms in the same dataset, field inheritance has priority over dataset inheritance.

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 every 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](#)

Dataset inheritance

Dataset inheritance declaration

The dataset inheritance mechanism is declared from DMA in the data model (Configuration > Data model properties) or directly in the XML Schema Definition, when using (Actions > Import 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:

<i>root record</i>	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.
<i>overwriting record</i>	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.
<i>inherited record</i>	Not locally defined in the current table and has a parent record. All values are inherited. Functions are always resolved in the current record context and are not inherited.
<i>occulting record</i>	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](#)

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](#)

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 the node's advanced properties in the underlying data model.
- 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>

        <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.

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.

Permissions

Permissions dictate the access each user has to data and actions.

- [Overview](#)
 - [Users, roles and profiles](#)
 - [Permission rules](#)
 - [Resolution of permissions](#)
 - [Owner and administrator special permissions](#)
 - [Impact of merge on permissions](#)
- [Important considerations about permissions](#)
 - [Actions and user services granting high privileges](#)
 - [API access without permission checks](#)
 - [Using permission for hiding information in the UI](#)
 - [Limitations of the permission checks in Query API](#)
 - [Scripted permission rules on records and table history](#)
 - [Using hidden fields in custom display labels](#)
 - [Linked field permission check](#)
 - [Table action permission related limitations](#)
 - [Permission cache life cycle](#)
- [Defining confidentiality](#)
- [Defining user-defined rules](#)
 - [Defining dataspace user-defined rules](#)
 - [Defining dataset user-defined rules](#)

- Defining dynamic rules
 - Defining scripted permission rules on data
 - Defining access rules on data
 - Defining activation rules on service
 - Defining permission rules on service
- Resolving permissions on data
 - Resolving user-defined rules
 - Resolving dynamic rules
- Resolving permissions on services
 - Resolving user-defined rules
- Resolving permissions on actions
 - Resolving user-defined rules

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](#).

Special definitions:

- A *built-in 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](#).

Dynamic permission rules can be either programmatic rules created by developers, or scripted rules created by administrators. See the section [Defining dynamic rules](#).

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.

Dynamic permissions are always considered to be restrictive.

Note:

In the Java API, the class `SessionPermissions` provides access to the resolved permissions.

See also:

- [Resolving permissions on data](#)

- Resolving permissions on services
- Resolving permissions on actions

Owner and administrator special permissions

On a dataset

A built-in 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 a built-in 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.

A built-in 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 a built-in 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

In this section are listed some very important information that must be kept in mind while working with permissions.

Actions and user services granting high privileges

The following actions and their related user services must only be allowed to trusted administrators:

- 'Create dataspace' (gives the 'owner' role, which grants the right to define the dataspace permissions)
- 'Create dataset' (gives the 'owner' role, which grants the right to define the dataset permissions)
- 'Import archive' (allows writing the archive content regardless of any permission)

Note:

See the [Owner and administrator special permissions](#) section for more information about the privileges granted to these profiles.

API access without permission checks

Developers and administrators must be aware that some parts of the API can run without any permission check. In general if the code run in a context with a `Session` provided, it means that permissions will be checked. Here are some specific cases where permissions are not checked:

- When a Java procedure disables all permission checks by using `ProcedureContext.setAllPrivileges`.
- When accessing EBX data by directly querying your database, in the case a table enables the [replication mode](#), or the [historization](#). This is because EBX permissions are not "translated" in the underlying database. As a consequence, either the database access must be globally restricted or proper permissions must be defined in it.

Using permission for hiding information in the UI

Using the permissions only to hide in the UI some non sensitive information is unadvised, especially if this information is likely to be used for filtering / joining / sorting in some queries. In such cases, UI-only hiding methods should be used instead. For instance by setting the field as [hidden for default views](#) in the datamodel property and/or by creating [views](#) for the concerned users.

When the above suggestions can not be applied and permissions have to be used to hide non sensitive information, the nodes must be set as non confidential in the model. For more information, see [Defining confidentiality](#).

Limitations of the permission checks in Query API

The permission check performed when specifying a session in a `Query` or `Request` will throw a `QueryPermissionException` if any confidential field used in the query is hidden for the current user. All fields are considered confidential except:

- Fields belonging to primary keys;
- Fields explicitly defined as non-confidential. For more information, see [Defining confidentiality](#);

Attention:

Confidential nodes with a record-dependent `AccessRule` are only partially confidential. Indeed, the method `AccessRule.getPermission` is only called with a dataset as the

`aDataSetOrRecord` parameter while performing the permission check, never with a record. As a consequence, the record-dependent logic of the rule will not apply during this check.

Scripted permission rules on records and table history

There is currently a limitation when a table has both the history activated and a scripted permission rule on record specified: for security reason access to the table history is totally disabled for everyone. Access to history will be allowed in a future version.

Using hidden fields in custom display labels

Resolution of custom display labels for tables ('defaultLabel' property) and relationships ('display' property) takes into account permission. As soon as an hidden field is detected in the label, the primary key will be displayed instead.

Note:

This is not the case when using API like `TableRefDisplay` or `Adaptation.getLabelOrName`. Since the provided contexts do not contain the current session, no permission check can be performed. As a consequence, developer should make sure that no confidential data is exposed when using these APIs.

Note:

Also note that quick search will ignore nodes with hidden fields in custom display label in the context of history view and/or in a child dataset.

Because of this behavior it is highly discouraged to use labels for filtering in a query. When labels with hidden fields are used, it will be replaced by the pk value and the filter will become inconsistent.

Linked field permission check

When a [linked field](#) access permission is computed, the result is the minimum between the permission applying to the node in the main table and the node in the target table. Practically it

means that if a field is hidden in a table, all linked fields pointing on it in other tables will also be hidden.

Table action permission related limitations

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.getNodeAccessPermission` beforehand in the procedure.

Permission cache life cycle

To optimize the resolution of permissions for both data and user services, a dedicated cache is implemented at the session level. All permissions are cached including dynamic rules, it means that a rule result should not change for the duration of the cache which is explained below.

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.clearCache`. Otherwise, these modifications will not be reflected until the end of the procedure.

Defining confidentiality

The confidentiality of a node defines whether or not it can be used in a `Query` or `Request` when hidden for a specific user. By default, all nodes are confidential.

Attention:

It is recommended to be very vigilant before setting a field as non confidential because although hidden, a malicious user could then "guess" the value by filtering or sorting using this node. It can not always be done, but it is generally a better option to use views to handle these kinds of use cases, as recommended in [Using permission for hiding information in the UI](#).

The confidentiality of nodes can be configured in the DMA. For more details, see [Node confidentiality](#).

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	• Can view the dataspace. • Can access datasets according to dataset permissions.
Read-only	• Can view the dataspace and its snapshots. • Can view child dataspace, if allowed by permissions. • Can view contents of the dataspace, though cannot modify them.
Hidden	• Can neither see the dataspace nor its snapshots. • If allowed to view child dataspace, can see the current dataspace but cannot select it.

DATASPACE ACCESS	AUTHORIZATION
	• Cannot access the dataspace contents, including datasets. • Cannot perform any actions on the dataspace.
Restriction policy	Indicates whether this dataspace profile-permission association should have priority over other permissions rules.
Create a child dataspace	Indicates whether the profile can create child dataspace from the current dataspace.
Create a child snapshot	Indicates whether the profile can create snapshots of the current dataspace.
Initiate merge	Indicates whether the profile can merge the current dataspace with its parent dataspace.
Export archive	Indicates whether the profile can export the current dataspace as an archive.
Import archive	Indicates whether the profile can import an archive into the current dataspace.
Close a dataspace	Indicates whether the profile can close the current dataspace.
Close a snapshot	Indicates whether the profile can close a snapshot of the current dataspace.
Rights on services	Indicates if a profile has the right to execute services on the dataspace. By default, all dataspace services are allowed. A built-in 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 as usual.

Defining dataset user-defined rules

For a given dataset, the allowable permissions for each profile are as follows:

Actions on datasets

Restriction policy	Indicates whether this dataset profile-permission association should have priority over other permissions rules.
Create a child dataset	Indicates whether the profile has the right to create a child dataset of the current dataset.
Duplicate dataset	Indicates whether the profile has the right to duplicate the current dataset.
Change the dataset parent	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:

Create a new record	Indicates whether the profile has the right to create records in the table.
Overwrite inherited record	Indicates whether the profile has the right to overwrite inherited records in the table.
Occult inherited record	Indicates whether the profile has the right to occult inherited records in the table.
Delete a record	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.

Read-write	Can view and modify node values.
Read	Can view nodes, but cannot modify their values.
Hidden	Cannot view nodes.

Permissions on services

A built-in administrator or an owner of the current dataspace can modify the service default permission to either restrict or grant access to certain profiles.

Enabled	Grants service access to the current profile.
Disabled	Forbids service access to the current profile. It will not be displayed in menus, nor will it be launchable via web components.
Default	Sets the service permission to enabled or disabled, according to the default permission defined upon service declaration.

See

`ActivationContext.setDefaultPermission`

for more information.

Defining dynamic rules

Dynamic 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 `AccessRule`, described in the section below [Defining access rules on data](#).
- the scripted record permission rule, described in the section below [Defining scripted permission rules on data](#).
- the `ServiceActivationRule`, described in the section below [Defining activation rules on service](#).
- the `ServicePermissionRule`, described in the section below [Defining permission rules on service](#).

Defining scripted permission rules on data

scripted permission rules are rules that dynamically define, depending on the context, the read/write rights on the records of a table.

To define such a rule, a [record permission script](#) must be created in the DMA. A script editor is available on the table node definition, in the "Extensions" tab.

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 `AccessRule` or `AccessRuleForCreate` interface.
2. Assignment of this rule to concerned nodes in the schema extension: `SchemaExtensions`.

According to the rule target (model node(s) or records) and type (`AccessRule` or `AccessRuleForCreate`), several methods such as

`SchemaExtensionsContext.setAccessRuleOnOccurrence` or `SchemaExtensionsContext.setAccessRuleForCreateOnNode` 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](#) 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 `ServiceActivationRuleForDataspace` interface or `ServiceActivationRuleForDataset`, 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.setActivationRule` or `ActivationContextWithDatasetSet.setActivationRule` methods.

The resulting assigned rule will be evaluated during the service activation evaluation. See [Resolving permissions on services](#) 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` interface.
2. Assignment of this rule to the impacted services:
 - Either, for new services, at their declaration level via the `ActivationContext.setPermissionRule` 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](#) for more information.

- Or, for existing services, in the [schema extension](#) via the `SchemaExtensionsContext.setServicePermissionRuleOnNode` and `SchemaExtensionsContext.setServicePermissionRuleOnNodeAndAllDescendants` 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](#) 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.

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
User 1	• user1 • role A • role B
User 2	• user2 • role A • role B • role C
User 3	• user3 • role A • role C

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

PROFILE	ACCESS RIGHTS	RESTRICTION POLICY
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
User 1	Hidden
User 2	Read
User 3	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 a built-in 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](#).

The final resolved access rights will be the minimum between the dataspace, dataset and the resolved access right for *N*.

Resolving dynamic rules

There are three levels of resolution for dynamic 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. However, a scripted rule can be specified on top of a programmatic rule at the table level.

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.setAccessRuleOnOccurrence` 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.setAccessRuleOnNode` 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 dynamic rules, if any.

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](#).

See also:

- [Overview](#)

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.defineActivation` method;
- the potential dynamic activation rule ([ServiceActivationRule](#)) also defined within the `UserServiceDeclaration.defineActivation` 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](#) section.

- If a [global permission rule](#) is defined for the service, it must return `enabled` for the context provided (see `ServicePermissionRuleContext`).
- If a [local permission rule](#) is defined for the selected node, it must return `enabled` for the context provided (see `ServicePermissionRuleContext`).

Resolving user-defined rules

Example

In this example, there are two users belonging to different roles and profiles:

User	Profiles

User 1	• user1 • role A • role B
User 2	• role C • role D

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](#)

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](#)) 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
Dataspace	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	
Dataset	Duplicate the dataset
Delete the dataset	
Activate/deactivate the dataset	
Create a view	
Table	Create a new record

ACTION OBJECT	AVAILABLE ACTIONS
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](#) section.

Resolving user-defined rules

Example

In this example, we have two users belonging to different roles and profiles:

User	Profiles
User 1	• user1 • role A • role B
User 2	• role C • role D

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

PROFILE	CREATE A RECORD	OVERRIDE A RECORD	OCCULT A RECORD	DELETE A RECORD	RESTRICTION POLICY
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](#)

Criteria editor

- [Overview](#)
- [Conditional blocks](#)
- [Atomic criteria](#)
 - [Expression](#)

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](#)

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.

Atomic criteria

An atomic predicate is defined by a field, an operator, and an expression (either a value or an XPath formula).

Field	Specifies the field of the table to which the criterion applies.
Operator	Specifies the operator used. Available operators depend on the data type of the field.
Value	Specifies the value or expression. See Expression below.
Code only	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.

Search

- [Overview](#)
 - [Value-labeling](#)
 - [Limitations](#)
- [Search strategies for string fields](#)
 - [Basic built-in strategies for strings](#)
 - [Advanced strategies for strings](#)
 - [Default strategy for string fields](#)
 - [Advanced custom strategies](#)
- [Search strategy for primary key fields](#)
- [Search strategy for foreign key fields](#)
- [Excluding a field from search \('Void' indexing\)](#)
- [Assigning a search strategy to a field](#)

Overview

A search strategy defines how a field is indexed and queried. Any field is associated with a default search strategy, primarily based on its data type.

Search strategies are specified in the Data Model Assistant:

- when editing a field, its search strategies can be set in the 'Extensions' tab;
- at the data model level, custom search strategies can be specified, under 'Extensions > Search' element in the left pane;

See also:

- [Quick Search](#)

Value-labeling

Value-labeling is a global feature in EBX® to display user-friendly labels instead of raw values. For example, in the user interface, a foreign key field displays the label of the linked record, or a field based on a static enumeration displays the localized label associated with the raw value, as specified by the data model.

If a field supports value-labeling, the Quick search and the sort in the user interface usually apply on the displayed label, to preserve an intuitive user interface.

There are some exceptions, where raw value is still used by the quick search and the sort operation:

- Programmatic labels and programmatic enumeration constraints (a foreign key specifying a `TableRefDisplay` or whose display depends on a `UILabelRenderer` specified on the target table, or a field constrained by a `ConstraintEnumeration`). It is recommended to use alternative solutions (display patterns and foreign keys).
- Enumeration constraint defined using another node (`<osd:enumeration osd:path=...`). It is recommended to use an alternative solution (a foreign key).

Obviously, if a field is displayed through a `UIWidget` (or a `UIBean`), to preserve an intuitive user interface, it is expected for the custom component to display the label (or the value, if this field does not enable value-labeling).

Limitations

In general, the following fields are not included in the Quick search and they are not optimized for other operators:

- computed fields with non-local dependency;
- inherited fields.

In the specific cases of inherited dataset, history view or mapped tables, legacy search is used. This implies that the size of the table cannot be quickly estimated, and might not be presented in the UI. It also implies that Quick search:

- considers all searchable fields (including computed fields with non-local dependency);
- behaves like a 'contains' (Lucene syntax cannot be used);
- does not support sort by relevancy;
- may perform poorly on tables with large volumes.

Regarding the Advanced Search pane, all fields will be available, except those of type `osd:locale` which are not defined as enumerations, and those of type `osd:resource`.

In the case of a node defining a display pattern, non sortable search strategies are prohibited for the default search template. Those are yet allowed in the other search templates.

Search strategies for string fields

Basic built-in strategies for strings

'Text'	The 'Text' search strategy is intended to contain multiple words, such as descriptions, texts or comments. This strategy supports full-text search and fuzzy search. Sorting, and some functions such as the 'equals' and 'starts-with' operators, are irrelevant and are not supported. This strategy is lightweight, consuming little disk space. See also Quick Search
'Code'	The 'Code' search strategy is intended for codes and identifiers. Values are considered as one single token, allowing any kind of case-sensitive and case-insensitive filter. Full-text search is irrelevant and is replaced by a 'contains' by default. This can be modified to a 'starts-with' by defining a custom 'Code' strategy. For large volumes, a starts-with is preferable to achieve better performance.
'Name'	The 'Name' search strategy is intended for names and labels that contain only a few words. Besides having the same search capabilities as 'Text', 'Name' strategy also allows sort, and supports the same filters as 'Code'. This strategy has the most capabilities, but consumes more disk space.

If the purpose of the field allows it, it is advised to choose the 'Text', 'Code' or 'Excluded from search' strategy, rather than this one.

Advanced strategies for strings

Advanced search strategies are meant to support searching using alternative fuzzy algorithms, and does not support sorting or filtering operations. It is advised to use them in a custom search template, and use a basic strategy with the default search template.

'Levenshtein'	The 'Levenshtein' strategy is intended for names and labels that contain a few words. This strategy only supports fuzzy search, based on an edit distance algorithm. The search syntax (+,-,...) is not supported. This strategy does not perform well on large volumes. The parameters of the strategy allow favorizing the performance, by setting the maximum edit distance to 1, and using an invariant prefix.
'Soundex'	The 'Soundex' strategy is intended for names and labels that contain a few words. This strategy only supports fuzzy search, based on the Soundex approximate phonetic algorithm. This algorithm is meant for english, and ignores digits.
'Double Metaphone'	The 'Double metaphone' strategy is intended for names and labels that contain a few words. This strategy only supports fuzzy search, based on the Double metaphone approximate phonetic algorithm. This algorithm is more modern than Soundex, and gives better results, especially for non-english languages. Digits are ignored.

'NGram strategy'

The 'NGram' strategy is intended for names and labels that contain a few words. This strategy only supports fuzzy search, based on a distance algorithm. This algorithm splits the values into smaller sequences called 'grams'. It performs better than the 'Levenshtein' search strategy, but consumes more disk space.

Default strategy for string fields

The 'Name' strategy is applied to string fields by default, except:

- If the field is part of the primary key, it is set by default to 'Code'.
- If the field is a foreign key, it is forced to 'Code' and cannot be changed.
- If the field has a built-in datatype extending `xs:string`, then it has a strategy relevant to its datatype; for instance `osd:text`, `xs:Name`, `osd:email`, `osd:html`, etc.

As the default strategy 'Name' can be irrelevant and consumes more disk space, the data model compilation reports warnings for fields with the 'Name' strategy set as default, so as to ensure that strategies are defined on purpose. We advise to choose the 'Text' strategy, when the length of the expected values is greater than 80, as a rough estimate. Long values (> 32766 bytes once encoded into UTF-8) will not be fully indexed with the 'Name' or 'Code' strategy. Quick search is not affected, but sorting will consider only the first 1000 characters, and some operators ('equals' and 'ends-with', SQL DISTINCT and COUNT DISTINCT, ...) will not return the correct results.

Advanced custom strategies

Some strategies accept parameters, for example to define stop words, or a specific language. This is done by creating a record in the 'Search strategies' table of the 'Search' data model extension. The new parameterized strategy will be available for selection in the 'Extension' tab, for compatible fields.

Search strategy for primary key fields

Primary key fields must have a sortable search strategy defined on the default search template. This excludes the 'Void' strategy for all data types, and the 'Text' strategy for strings. Do note that it is still possible to use non sortable search strategy for a primary key field, if it is defined within a search template other than the default search template.

Search strategy for foreign key fields

Foreign key fields have two levels of search:

- First, if applicable, the search is performed on each field of the displayed label of the foreign key. Each field strategy is inherited from the field in its target table. This first level is not always applicable, for instance when the search string cannot be converted to any of the target field data types.
- Secondly, when the first level of search cannot be applied, the search is performed on the string representation of the target primary key. Modifying the search strategy of a foreign key field in the 'Extension' tab in the Data Model Assistant only affects this second level of search. It can only be a 'Code' search strategy (built-in or customized).

See also:

- `PrimaryKey.syntax`

See also:

- [Value-labeling](#)

Excluding a field from search ('Void' indexing)

The 'Excluded from search' (or `Void`) strategy deactivates indexing, making filter, search, or sort impossible. It is available for all data types, and is intended for fields that are never queried. Values can still be accessed through their record. Disabling the indexing reduces the disk space consumed and speeds up some operations like data import.

Assigning a search strategy to a field

A search strategy can be associated with a field, by means of a [search template](#). This is done in the 'Extension' tab of the field, in the Data Model Assistant. Assigning multiple search strategies to a field requires registering additional search templates into a module. Only the addons EBX® Information Search and EBX® Match and merge are concerned by additional search templates.

Performance and tuning

- Environment
 - Memory management
 - Disk
 - CPU
 - Using the native LZ4 library
 - Scanning on server startup
- Database
 - Reorganizing database tables
- Data modeling
 - Aggregated lists
- Data validation
- Accessing tables
 - Functionalities
 - Query on tables
 - REST access to history table
 - Setting a fetch size
- Performance checklist for other Java customizations
 - Expensive programmatic extensions
 - Unnecessary index refresh
 - Transaction threshold for mass updates
 - Triggers
 - Directory integration

Environment

Memory management

Memory allocated to the application server	Since the query engine retrieves the necessary information from persistent storage, the memory allocated to the Java Virtual Machine (usually specified by the -Xmx parameter) can be kept low. We recommend setting this figure to less than 1/3rd of the total memory available on the OS. We also advise to stay below 32 GB, which should fit all reasonable use cases, and allow benefiting from the Compressed Ordinary Object Pointer feature.
Memory allocated to the operating system	On the OS running the application server, it is important to leave sufficient room to the OS cache, letting it optimize access to the persistent Lucene indexes. Indeed, once these have been loaded from the file system, the OS uses its memory cache to speed up subsequent accesses to this same data, and avoid reloading it every time from the disk. This is only possible if sufficient RAM has been left for this purpose. It is also necessary to configure the OS so that the JVM process can reserve the resources required by numerous memory-mapped files (see this article for details). On a Linux OS, this can be done by issuing the following commands: <pre>ulimit -n 512000</pre> <pre>sysctl vm.max_map_count=262144</pre>

Memory monitoring	Indications of EBX® load activity are provided by monitoring the underlying database, and also by the 'monitoring' logging category. If the numbers for <i>cleared</i> and <i>built</i> objects remain high for a long time, this is an indication that EBX® is swapping on the application server. In that case, the memory allocated to the application server should be increased.
Garbage collector	Tuning the garbage collector can also benefit overall performance. This tuning should be adapted to the use case and specific Java Runtime Environment used.

Disk

The EBX® repository data are indexed into Lucene indexes, stored on the disk under the [root directory](#).

Disk space: a rule of thumb for the disk size is to plan for 10 times the space occupied by the table `G_BLK` in the relational database.

Disk latency: in order to maintain good overall performance, it is important for the disk storing the Lucene indexes to have low latency.

See also:

- [Setting the EBX® root directory](#)

CPU

The number of CPUs available for the application server must be defined considering the number of concurrent HTTP requests to be served, the complexity (CPU cost) of the implied tasks, and the background activities, including the Java garbage collection.

Large imports and more generally large transactions involving many creates, updates and deletes will be completed faster if the difference between the server load and the number of available processors allows the indexing to be efficiently run in parallel within the transaction. The `persistence` log category contains the following entries:

- *'Setting forced synchronous indexing to false (...)'* to indicate that the indexing will be performed concurrently;
- *'Setting forced synchronous indexing to true (...)'* to indicate that the indexing will not be performed concurrently.

The difference mentioned above is assessed every ten seconds and is computed using the methods `getSystemLoadAverage()` and `getAvailableProcessors()` in the Java class `java.lang.management.OperatingSystemMXBean`. Both numbers are written at the end of the log entry above.

Using the native LZ4 library

The LZ4 library is used to store data to and retrieve data from the database. To speed up data access, it is required to perform a `ebx-lz4.jar` native installation.

See [Data compression library](#) for more information.

Scanning on server startup

To speed up the web applications server startup, the [JAR files scanner](#) should be configured.

Database

Reorganizing 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](#).

A specificity of EBX® is that creating dataspace and snapshots adds new entries to tables `GRS_DTR` and `GRS_SHR`. 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

Data modeling

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](#), 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](#) 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.

Data validation

The internal validation framework will optimize the work required during successive requests to update the validation report of a dataset or a table. The incremental validation process behaves as follows:

- The first call to a dataset or table validation report performs a full validation of the dataset or the table.
- The next call to the validation report will compute the changes performed since the last validation. The validation report will be updated according to these changes.
- Validation reports are stored persistently in the TIBCO EBX® repository. This reduces the amount of memory dedicated to validation reports when datasets have a large amount of validation messages. Also, validation reports are not lost when the application server restarts.

- Validation reports can be reset manually in the user interface by an administrator user (this option is available from the validation report section in EBX®). As a consequence, resetting validation reports must be used with caution since associated datasets or tables will be fully revalidated during the next call to their validation reports.

See `Adaptation.resetValidationReport` for more information.

Certain constraints are systematically re-validated, even if no updates have occurred since the last validation. These are the constraints with *unknown dependencies*. An element has unknown dependencies if:

- It specifies a [programmatic constraint](#) in the default *unknown dependencies* mode.
- It declares a [computed value](#), or it declares a dynamic facet that depends on an element that is itself a [computed value](#).
- It is an [inherited field](#) or it declares a dynamic facet that depends on a node that is itself an [inherited field](#).

Consequently, on large tables, it is recommended to:

- Avoid constraints with unknown dependencies (or at least to minimize the number of such constraints). 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](#).
- To use [constraints on tables](#) instead of [programmatic constraints](#) defined at field level. Indeed, if a table defines constraints at field level, then the validation process will iterate over all the records to check if the value of the associated field complies with the constraint. Using [constraints on tables](#) gives the opportunity to execute optimized queries on the whole table. See [below](#) for recommendations on how to optimize such queries. However, there is a trade-off: A [constraint on tables](#) does not benefit from the incremental validation and will always be validated completely. If in general only a few records are created or deleted between validations, a [programmatic constraint](#) might be more performant. Benchmarks or a sampling profiler can help to make the decision.
- Avoid the use of the facet [pattern](#) since its check is not optimized on large tables. That is, if a field defines this facet then the validation process will iterate over all the records to check if the value of the associated field complies with the specified pattern.

The following properties can be used to minimize the impact on performance when logging a validation report:

- Set the `ebx.validation.report.logContent` property to `false` to avoid logging the individual validation messages. The validation report summary is still logged, including the message count. You can also avoid logging the validations started by the Java API completely. See `ValidationSpec.setResultLogged` for more information.
- Though less effective than the previous recommendation, setting the `ebx.validation.report.maxItemDisplayed` property to a lower value than the default 100 reduces the amount of validation messages to log, thereby reducing the related computational work.

Accessing tables

Functionalities

Tables are commonly accessed through EBX® UI, data services and also through the `Request` and `Query` APIs. 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.getValue`.
- **Filtering:** An [XPath predicate](#), a [programmatic filter](#), or a record-level [permission rule](#) requires a selection of records.
- **Sort:** A sort of the resulting records can be performed.

Query on tables

Architecture and design

In order to improve the speed of operations on tables, persistent Lucene indexes are managed by the EBX® engine.

Attention:

Faster access to tables is ensured if indexes are ready and maintained in the OS memory cache. As mentioned [above](#), it is important for the OS to have enough space allocated.

Performance considerations

The query optimizer favors the use of indexes when computing a request result. If a query cannot take advantage of the indexes, it will be resolved in Java memory, and experience poor performance on large volumes. The following guidelines apply:

Attention:

- Only XPath predicates and SQL queries can benefit from index optimization.
- Some fields and some datasets cannot be indexed, as described in section [Limitations](#).
- XPath predicates using the `osd:label` function cannot benefit from index optimization

If indexes have not yet been built, additional time is required to build and persist the indexes, on the first access to the table.

Accessing the table data blocks is required when the query cannot be computed against any 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.

Apart from that, there are other considerations that can impact requests / queries performance:

- If the same request / query is going to be executed multiple times with just a variation on the filter value, it is recommendable to use parameters in said filter:
 - `request.setXPathFilter("./field=$param");`
 - `SELECT * FROM myTable WHERE field=?`
- If possible, use `RequestResult.isSizeGreaterOrEqual` rather than `RequestResult.getSize`
- If possible, avoid `RequestResult` size-related calls. Especially avoid the unnecessary code pattern "check isEmpty before traversing" on a `RequestResult`:

```

try (RequestResult result = ...) {

    if (result.isEmpty())

        return; // unnecessary and inefficient!

    for (Adaptation record : result) {

        ...

    }

}

// simply do:

try (RequestResult result = ...) {

    for (Adaptation record : result) {

        ...

    }

}

```

It is possible to get information through the [memory monitoring](#) and request logging categories.

Accessing and modifying a table

The following access lead to poor performance, and must be avoided:

- Access a table after a few modifications, repeatedly. It implies the index state to be refreshed after each modification. The cost of refreshing makes this pattern ineffective. Instead, perform a single query and apply the modification when browsing the results.
- If there is an ongoing access to the same table, concurrently to the previous case, it prevents outdated index files to be deleted. As a consequence, the size of the index on disk increases, and the server may run out of disk space in extreme cases. When the concurrent access is

closed, the index size is back to normal. This is usually a sign that a Request or a Query is not properly closed.

See also:

- `RequestResult.close`
- `QueryResult.close`

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` parameter must be set to `false`.

See [History](#) for more information.

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.

- On a history table, 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`
- `RequestResult`

Performance checklist for other Java customizations

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	• programmatic constraints • computed values
Table access	• record-level permission rules • programmatic filters
EBX® content display	• computed values • UI Components • node-level permission rules
Data update	• triggers

For large volumes of data, using algorithms of high computational complexity has a serious impact on performance. For example, the complexity of a constraint's algorithm 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 by code analysis or by using a Java profiling tool.

Unnecessary index refresh

Refreshing a Lucene index takes time. It should be avoided whenever possible.

When does a refresh happen?	In the context of a transaction, an index refresh occurs when the table has been modified and one of the conditions below occurs: 1. For a lookup by primary key, the refresh is always triggered if the searched key has been "touched" (created, modified or deleted) in the current <code>Procedure</code> (or <code>TableTrigger</code>). 2. For a standard <code>Query</code> (or <code>Request</code>), an index refresh is always performed if the table has been modified in the current <code>Procedure</code> (or <code>TableTrigger</code>).
Coding recommendations	1. To avoid triggering a refresh through a lookup by primary key, the developer must register the <code>Adaptation</code> object returned from the last call to <code>doCreateOccurrence</code> or <code>doModifyContent</code> , and reuse this object instead of performing the lookup.

2. Avoid any lookup by primary key on a record that has been deleted in the current procedure.
3. In the case of a query triggering the refresh, the developer must ask the following question: can this query be avoided in my procedure?

Transaction threshold for mass updates

It is generally not advised to use a single transaction when the number of atomic updates in the transaction is beyond the order of 10⁵. Large transactions require a lot of resources, in particular, memory, from EBX® and from the underlying database.

To reduce the transaction size, it is possible to:

- Specify the property `ebx.manager.import.commit.threshold`. However, this property is only used for interactive archive imports performed from the EBX® user interface.
- Explicitly specify a `commit threshold` inside the batch procedure.
- Structurally limit the transaction scope by implementing `Procedure` 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.setTriggerActivation
```

.

Directory integration

Authentication and permissions management involve the [user and roles directory](#).

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.isUserInRole`.

Administration Guide

- [Administration overview](#)
- [Installation & configuration](#)
 - [Supported environments](#)
 - [Java EE deployment](#)
 - [Installation notes](#)
 - [Installation note for JBoss EAP 7.4.x](#)
 - [Installation note for WebSphere Application Server Liberty 23.x](#)
 - [Installation note for Tomcat 9.x](#)
 - [Installation note for WebLogic 14c](#)
 - [TIBCO EBX® main configuration file](#)
 - [Initialization and first-launch assistant](#)
 - [Deploying and registering TIBCO EBX® add-ons](#)
- [EBX® Container Edition](#)
 - [Building the image](#)
 - [Running the image](#)
 - [Customizing the image](#)
- [Technical administration](#)
 - [Repository administration](#)
 - [UI administration](#)
 - [UI – Workflow launcher](#)
 - [Users and roles directory](#)
 - [Data model administration](#)
 - [Database mapping administration](#)
 - [Workflow management](#)
 - [Task scheduler](#)

- [Audit trail](#)
- [Collaborative features \(Team Up\)](#)
- [Other](#)
- [Distributed Data Delivery \(D3\)](#)
 - [Introduction to D3](#)
 - [D3 broadcasts and delivery dataspace](#)
 - [D3 JMS Configuration](#)
 - [D3 administration](#)
- [Staging](#)

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.

- [Repository management](#)
 - [Object cache](#)
 - [Obsolete contents](#)
 - [Workflow](#)
- [Disk space management](#)
 - [Purge of logs](#)
 - [XML audit trail \(deprecated\)](#)
- [Data model](#)
 - [Publication management](#)
 - [Refresh data models](#)
- [Data validation](#)
 - [Resetting validation reports](#)
- [Perspectives](#)

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](#) and [Deleting dataspaces, snapshots, and history](#).

It is also possible to archive files of the file system type in order to reduce the storage costs, see [EBX® monitoring](#).

Administration tasks can be scheduled by means of the task scheduler, using built-in tasks, see [Task scheduler](#).

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 report.

See also:

- [Memory management](#)

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](#).

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](#) 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](#) for more information.

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 size (in megabytes) 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
```

XML audit trail (deprecated)

EBX® is provided with an XML audit trail manager, now disabled by default. Any customized management (including purge, backups, etc.) is the user's responsibility.

See [Activating the legacy XML audit trail \(deprecated\)](#) and [Legacy XML Audit trail \(deprecated\)](#) for more information.

Data model

Publication management

The management of publications of [embedded data models](#). See [Data model administration](#) 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](#) for more information.

Data validation

Resetting validation reports

Validation reports can be reset from the administration menu. It is possible to reset the validation reports of all the datasets contained in the repository or only those of the datasets selected in the tabular view. Therefore, the related datasets will be fully validated at the next explicit validation request.

See `Adaptation.resetValidationReport` for more information.

Perspectives

EBX® offers extensive UI customization options. Simplified interfaces ([Recommended perspectives](#)) 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](#) for more information.

Installation & configuration

- [Supported environments](#)
- [Java EE deployment](#)
- [Installation notes](#)
 - [Installation note for JBoss EAP 7.4.x](#)
 - [Installation note for WebSphere Application Server Liberty 23.x](#)
 - [Installation note for Tomcat 9.x](#)
 - [Installation note for WebLogic 14c](#)
- [TIBCO EBX® main configuration file](#)
- [Initialization and first-launch assistant](#)
- [Deploying and registering TIBCO EBX® add-ons](#)

Supported environments

- [Browsing environment](#)

- Supported web browsers
- Screen resolution
- Refreshing pages
- 'Previous' and 'Next' buttons
- Zoom troubleshooting
- Browser configuration
- Limitations
- Supported application servers
- Supported databases

Browsing environment

Supported web browsers

The TIBCO EBX® web interface supports the following browsers:

Microsoft Edge Chromium	As Microsoft Edge Chromium 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.
Mozilla Firefox ESR 102 (see details)	As Mozilla Firefox is updated frequently, Cloud Software Group, Inc. only fully supports version ESR 102. See Mozilla Firefox ESR for more details.
Google Chrome	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®.

Limitations

Some browsers may have a limitation on the number of iframes that can be embedded. If this is the case, it limits to the number of items that can be pushed in the breadcrumb. Please check the browser documentation for more details.

Supported application servers

EBX® supports the following configurations:

- Java Runtime Environment: JRE 11 or 17.

Note: We advise to upgrade the runtime to a recent LTS version, to take advantage of the performance improvements it offers.

- Any Servlet/JSP container that complies with Servlet 3.0 (inclusive) up to 5.0 (exclusive): for example Tomcat 7.0 (inclusive) up to 10.0 (exclusive). Any Java Application Server supporting Java SE 11 and not using Jakarta EE 9 or greater: for example WebSphere Liberty 20 or higher for Java EE 8, WebLogic Application Server 14c or higher, JBoss EAP 7.4 or higher. See [Java EE deployment overview](#).
- The application server must support the JSON Processing 1.1 ([JSR 374](#)), or allow the uses of the implementation embedded in the `ebx.jar` library. For example, Tomcat does not provide any library to support this specification (only the embedded one can be used), WebLogic Application Server 14c or higher supports this specification, WebSphere Application Server Liberty and JBoss EAP allow including or excluding the available libraries.
- 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](#) 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.

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.

Oracle Database 19c or higher.	The distinction of <code>null</code> values bears certain limitations. On simple <code>xs:string</code> elements, Oracle does not support the distinction between empty strings and <code>null</code> values. See Empty string management for more information. The user with which EBX® connects to the database requires the following privileges: • CREATE SESSION, • CREATE TABLE, • ALTER SESSION, • CREATE SEQUENCE, • A non-null quota on its default tablespace.
PostgreSQL 11 or higher.	The user with which EBX® connects to the database needs the CONNECT privilege on the database hosting the EBX® repository. Other than this, the default privileges on the public schema of this database are suitable. Also, see this limitation regarding the evolution of datamodels in mapped modes.
Amazon Aurora PostgreSQL 11.16 (compatible with PostgreSQL 11.16 or higher).	The comments in the above section for PostgreSQL apply.
Google Cloud SQL for PostgreSQL 11 or higher.	The comments in the above section for PostgreSQL apply.

SAP HANA Database 2.0 or Higher.	When using SAP HANA Database as the underlying database, certain schema evolutions are not supported. It is, for example, impossible to reduce the length of a column; this is a limitation of HANA, as mentioned in the SQL reference guide: "For row table, only increasing the size of VARCHAR and NVARCHAR type column is allowed."
Microsoft SQL Server 2014 or higher.	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 history tables. The default database collation can be specified when the database is created. Otherwise, the collation of 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 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 2.1.212 or higher.

H2 is not supported for production environments.

The default H2 database settings do not allow consistent reads when records are modified.

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).

See also:

- [Repository administration](#)
- [Data source of the EBX® repository](#)
- [Configuring the EBX® repository](#)

Java EE deployment

- [Introduction](#)
- [Software components](#)
- [Embedded third-party libraries](#)
- [Required third-party libraries](#)
 - [Data compression library](#)
 - [Database drivers](#)
 - [SMTP and emails](#)
 - [Secure Socket Layer \(SSL\)](#)
 - [Java Message Service \(JMS\)](#)
- [Web applications](#)

- EBX® built-in web applications
- Custom web applications
- Deployment details
 - Introduction
 - User interface and web access
 - Data source of the EBX® repository
 - Mail sessions
 - JMS connection factory
 - JMS for data services
 - JAR files scanner
- Installation notes

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](#).

Software components

EBX® uses the following components:

- Library `ebx.jar`
- [Embedded](#) and [required](#) third-party Java libraries
- [EBX® built-in web applications](#) and optional [custom web applications](#)
- [EBX® main configuration file](#)
- [EBX® repository](#)
- [Default user and roles directory](#), integrated within the EBX® repository, or a third-party system (LDAP, RDBMS) for the user authentication

See also:

- [Supported environments](#)

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
- Apache Log4j 2 API
- ArcGIS API for JavaScript
- Javax Activation
- Javax Annotations
- Javax JSON Binding
- Javax SAAJ API
- Javax WS RS
- Javax XML Bind
- LZ4 compression for Java
- MicroProfile OpenAPI
- Simple Logging Facade for Java (SLF4J) API

For more information regarding the versions or the details of the Third-Party Library, please refer to the: `TIB_ebx_6.1.0_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.

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.

Data compression library

The library named `ebx-lz4.jar` must be deployed separately from `ebx.jar`. It contains several compression implementations: JNI dedicated architecture libraries and Java fallbacks. It is possible to ensure optimal compression and decompression performance for EBX® repository by following prerequisites. If prerequisites can not be validated, EBX® will function in Java fallbacks safe or unsafe, but its performance will be degraded. The default location for `ebx-lz4.jar` library is beside `ebx.jar`.

To verify the compression implementation actually used by the EBX® repository, please check the value of 'Compression' in 'Administration > System Information', section 'Repository information'. It should be '`JNI - validated`' for optimal performance. Otherwise, it will be '`Java[Safe | Unsafe] - validated`' for Java fallbacks.

Performance prerequisites

The JNI access is allowed to the following operating system architectures: `i386`, `x86`, `amd64`, `x86_64`, `aarch64` or `ppc64le`. To verify this value, please check the value of 'Operating system architecture' in 'Administration > System Information', section 'System information'.

To enable JNI access for `ebx-lz4.jar`, the library should be loaded by the **system class loader** (also known as the application class loader). The deployment may be done by following the [specific instructions for your application server](#).

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.

H2	Version 2.1.212 validated. Note that H2 is not supported in production environments. https://www.h2database.com/ If the H2 repository has been built with version 1.0 driver, this migration to v2 process must be applied.

Oracle JDBC	Oracle database 21c is validated on their latest patch set update. Determine the driver that should be used according to the database server version and the Java runtime environment version. Download the <code>ojdbc8.jar</code> certified library with JDK 8. Oracle database JDBC drivers download.
SQL Server JDBC	SQL Server 2012 SP4 and greater, with all corrective and maintenance patches applied, are validated. Remember to use an up-to-date JDBC driver, as some difficulties have been encountered with older versions. Include the <code>mssql-jdbc-8.4.1.jre8.jar</code> or <code>mssql-jdbc-8.4.1.jre11.jar</code> library, depending on the Java runtime environment version you use. Download Microsoft JDBC Driver 8.4.1 for SQL Server (zip).
PostgreSQL	PostgreSQL 10 and above validated Include the latest JDBC driver version 4.2 released for your database server and Java runtime environment. PostgreSQL JDBC drivers download.

See also:

- [Data source of the EBX® repository](#)

- [Configuring the EBX® repository](#)

SMTP and emails

According to the web application server being used, the library `JavaMail API` for email management may already be provided, or must be added manually.

EBX® requires a library that is compatible with version 1.5.6 of this API. See [Activating and configuring SMTP and emails](#) for more information on the configuration.

To facilitate manual installation, the `javax.mail-1.5.6.jar` has been provided and placed under the `ebx.software/lib/lib-mail` directory.

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>
- `ibmjsse.jar` : <https://www.ibm.com/developerworks/java/jdk/security/>

See also:

- [TIBCO EBX® main configuration file](#)

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 <https://www.oracle.com/java/technologies/java-ee-glance.html> 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 <https://www.oracle.com/java/technologies/java-se-glance.html> 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](#)

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](#) at the end of this chapter.

EBX® built-in web applications

EBX® includes the following built-in web applications.

WEB APPLICATION NAME	DESCRIPTION	REQUIRED
<code>ebx</code>	EBX® entry point, which handles the initialization on start up. See Deployment details for more information.	Yes
<code>ebx-root-1.0</code>	EBX® root web application. Any application that uses EBX® requires the root web application to be deployed.	Yes
<code>ebx-ui</code>	EBX® user interface web application.	Yes
<code>ebx-manager</code>	EBX® user interface web application.	Yes
<code>ebx-dma</code>	EBX® data model assistant, which helps with the creation of data models through the user interface. Note: The data model assistant requires the <code>ebx-manager</code> user interface web application to be deployed.	Yes
<code>ebx-dataservices</code>	EBX® data services web application. Data services allow external interactions with the EBX® repository using the SOAP operations and Web Services Description Language WSDL generation standards or using the Built-in RESTful services . Note: The EBX® web service generator requires the deployment of the <code>ebx-</code>	Yes

WEB APPLICATION NAME	DESCRIPTION	REQUIRED
	<code>manager</code> user interface web application.	

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](#)
- [Declaring modules as undeployed](#)

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](#)
- [Supported application servers](#)

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><load-on-startup>1</load-on-startup></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"/>
```

- When several EBX® Web Components are to be displayed on the same HTML page, for instance using iFrames, it may be required to disable the management of cookies due to limitations present in some Internet browsers.

For example, on Tomcat, this configuration is provided by the attribute `cookies` in the configuration file `server.xml` , as follows:

```
<Context path="/ebx" docBase="(...)" cookies="false"/>
```


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. In particular, the size of the connection pool must be set according to the number of concurrent users. See [Configuring the EBX® repository](#) 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> JBoss: <code>java:/EBX_REPOSITORY</code>	JDBC data source for EBX® Repository. Java type: <code>javax.sql.DataSource</code>

See also:

- [Configuring the EBX® repository](#)
- [Rules for the database access and user privileges](#)

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](#) 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
<code>mail/EBX_MAIL_SESSION</code>	Weblogic: <code>EBX_MAIL_SESSION</code> JBoss: <code>java:/EBX_MAIL_SESSION</code>	Java Mail session used to send emails from EBX®. 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](#) 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
<code>jms/</code> <code>EBX_JMSConnectionFactory</code>	Weblogic: <code>EBX_JMSConnectionFactory</code> JBoss: <code>java:/</code> <code>EBX_JMSConnectionFactory</code>	JMS connection factory used by EBX® to create connections with the JMS provider configured in the operational environment of the application server. Java type: <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](#) 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](#) 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](#) 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> JBoss: <code>java:/jms/EBX_QueueIn</code>	JMS queue for incoming SOAP requests sent to EBX® by other applications. Java type: <code>javax.jms.Queue</code>	No
<code>jms/EBX_QueueFailure</code>	Weblogic: <code>EBX_QueueFailure</code> JBoss: <code>java:/</code>	JMS queue for failures. It contains incoming SOAP requests for which	No

RESERVED RESOURCE NAME	DEFAULT JNDI NAME	DESCRIPTION	REQUIRED
	<code>jdbc/</code> <code>EBX_QueueFailure</code>	an error has occurred. This allows replaying these messages if necessary. Java type: <code>javax.jms.Queue</code> Note: 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 JMS in the EBX® main configuration properties for more information on these properties.	

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.jarsToSkip` property from the `catalina.properties` file.

Installation notes

EBX® can be deployed on any Java EE or Jakarta EE 8 application server that supports Servlet 3.0 up to 5.0 except. The following documentation on deployment and installation notes are available:

- [Installation note for JBoss EAP 7.4.x](#)

- Installation note for WebSphere Application Server Liberty 23.x
- Installation note for Tomcat 9.x
- Installation note for WebLogic 14c

Attention:

- The EBX® installation notes on Java EE or 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 custom or add-on 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](#) section.

Installation notes

- [Installation note for JBoss EAP 7.4.x](#)
- [Installation note for WebSphere Application Server Liberty 23.x](#)
- [Installation note for Tomcat 9.x](#)
- [Installation note for WebLogic 14c](#)

Installation note for JBoss EAP 7.4.x

- [Overview](#)
- [Requirements](#)
- [JBoss Application Server installation](#)
- [EBX® home directory configuration](#)
- [JBoss Application Server and Java Virtual Machine configuration](#)
- [JNDI entries configuration](#)
- [Data source and JDBC provider configuration](#)
- [EBX.ear application update](#)
- [EBX.ear application deployment](#)
- [EBX® application start](#)

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](#).
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar`, `ebx-lz4.jar` or other libraries in the class-loading system.

Requirements

- Java SE 11 or 17 ([update 7 and above](#))
- JBoss Application Server EAP 7.4
- Database and JDBC driver
- EBX® CD
- No CDI features in EBX®'s additional modules (since CDI will be automatically disable)

See also:

- [Supported environments](#)

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:
 - a. Select the language and click 'OK',
 - b. Accept the License and click 'Next',
 - c. Choose the installation path and click 'Next',

- d. Keep the 'Component Selection' as it is and click 'Next',
 - e. Enter 'Admin username', 'Admin password' and click 'Next',
 - f. On 'Installation Overview' click 'Next',
 - g. On 'Component Installation' click 'Next',
 - h. On 'Configure Runtime Environment' leave the selection as it is and click 'Next',
 - i. When 'Processing finished' appear, click 'Next',
 - j. Uncheck 'Create shortcuts in the start menu' and click 'Next',
 - k. Generate 'installation script and properties file' in the JBoss EAP 7.4 installation root directory,
 - l. 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>

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.

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.
3. Set the 'JBOSS_MODULES_SYSTEM_PKGS' environment variable like the following:

```
JBOSS_MODULES_SYSTEM_PKGS="org.jboss.byteman,net.jpountz"
```


4. Copy from the *EBX® CD*, the `ebx.software/lib/ebx-lz4.jar` Data compression library to a dedicated directory (for example `<JBOSS_HOME>/compress`).
5. Open the `standalone.sh` script file, placed in `<JBOSS_HOME>/bin` .
6. Create a 'CLASSPATH' environment variable like the following:

```
CLASSPATH="<path_to_the_data_compression_library>:${JBOSS_HOME}/jboss-modules.jar:${CLASSPATH}"
```

```
# For our example
```

```
# CLASSPATH="${JBOSS_HOME}/compress/ebx-lz4.jar:${JBOSS_HOME}/jboss-modules.jar:${CLASSPATH}"
```

7. Replace the launch command options for foreground and background executions like the following:

```
if [ "x$LAUNCH_JBOSS_IN_BACKGROUND" = "x" ]; then

    # Execute the JVM in the foreground

    eval `"$JAVA" -D"[Standalone]" $JAVA_OPTS \

        -cp "$CLASSPATH" \

        \ "-Dorg.jboss.boot.log.file="$JBOSS_LOG_DIR"/server.log" \

        \ "-Dlogging.configuration=file:"$JBOSS_CONFIG_DIR"/

logging.properties" \

        org.jboss.modules.Main \

        $MODULE_OPTS \

        -mp \"\"${JBOSS_MODULEPATH}\"\" \

        org.jboss.as.standalone \

        -Djboss.home.dir=\"\"$JBOSS_HOME\"\" \
```

```

-Djboss.server.base.dir=\\"$JBASS_BASE_DIR\" \

"$SERVER_OPTS"

JBASS_STATUS=$?

else

# Execute the JVM in the background

eval \\"$JAVA\" -D\"[Standalone]\" $JAVA_OPTS \

-cp \"$CLASSPATH\" \

\\-Dorg.jboss.boot.log.file=\"$JBASS_LOG_DIR\"/server.log\" \

\\-Dlogging.configuration=file:\"$JBASS_CONFIG_DIR\"/
logging.properties\" \

org.jboss.modules.Main \

$MODULE_OPTS \

-mp \\"$JBASS_MODULEPATH\" \

org.jboss.as.standalone \

-Djboss.home.dir=\\"$JBASS_HOME\" \

-Djboss.server.base.dir=\\"$JBASS_BASE_DIR\" \

"$SERVER_OPTS" "&"

...

fi

```

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"/>

```

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>
```

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](#) for more information.

SSL: see [Secure Socket Layer \(SSL\)](#) for more information.

JMS: see [Java Message Service \(JMS\)](#) 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.

EBX.ear application deployment

1. Copy EBX.ear into the `JBOSS_HOME/standalone/deployments` directory.

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](#) helps to configure the default properties of the initial repository.

Installation note for WebSphere Application Server Liberty 23.x

- [Overview](#)
- [Requirements](#)
- [WebSphere Application Server Liberty installation](#)
- [EBX® home directory configuration](#)
- [EBX® and third-party libraries deployment](#)
- [JNDI entries configuration](#)
- [Data source and JDBC provider configuration](#)
- [Java Virtual Machine configuration](#)
- [EBX® application deployment](#)

- [EBX® application start](#)

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](#).
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar`, `ebx-lz4.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.

Requirements

- Java SE 11 or higher. It's recommended to use [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](#)

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.

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.

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.
```

```
ebx-lz4.jar , coming from the ebx.software/lib/ directory of the EBX® CD.
```

The complete description of these components is given in the chapter [Java EE deployment](#). 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).

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](#).

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](#).

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"
```

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 create the EBX® library, edit the `<WLP_HOME>/usr/servers/<SERVER_NAME>/server.xml` file by adding the following lines into the `<server>` tag:

```
<library id="ebxLibLz4">

  <file name="${EBX_LIB}/ebx-lz4.jar"/>

</library>
```

Note:

WebSphere Application Server Liberty doesn't support `LZ4 JNI implementation`.

- 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="ebxApp" type="ear">

    <classloader commonLibraryRef="ebxLibLz4" />

</application>
```

EBX® application start

- 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/`
- At first launch, **EBX® Wizard** helps to configure the default properties of the initial repository.

Installation note for Tomcat 9.x

- [Overview](#)
- [Requirements](#)
- [Tomcat Application Server installation](#)
- [EBX® home directory configuration](#)
- [Tomcat Application Server and Java Virtual Machine configuration](#)
- [EBX® and third-party libraries deployment](#)
- [EBX® web applications deployment](#)
- [EBX® application start](#)

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](#).
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar`, `ebx-lz4.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.

Requirements

- Java SE 11 or 17
- Apache Tomcat 9.x
- Database and JDBC driver
- EBX® CD

See also:

- [Supported environments](#)

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

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.

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,\
```

```
ebx-lz4.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%

set CLASSPATH=<$CATALINA_HOME_or_$CATALINA_BASE>\compress\ebx-
lz4.jar;%CLASSPATH%
```

Where `<$CATALINA_HOME_or_$CATALINA_BASE>` must be replaced by

`%CATALINA_HOME%` or `%CATALINA_BASE%` if they have been configured. Otherwise this piece of text must be replaced by the Tomcat installation directory's path.

- 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="<$CATALINA_HOME_or_$CATALINA_BASE>\compress\ebx-lz4.jar;
%CLASSPATH%"
```

Where `<$CATALINA_HOME_or_$CATALINA_BASE>` must be replaced by

`%CATALINA_HOME%` or `%CATALINA_BASE%` if they have been configured. Otherwise this piece of text must be replaced by the Tomcat installation directory's path.

- 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="-Dorg.apache.catalina.connector.CoyoteAdapter.ALLOW_BACKSLASH=true $
ebx.properties"

export JAVA_OPTS="$EBX_OPTS
-Dorg.apache.catalina.connector.CoyoteAdapter.ALLOW_BACKSLASH=true $
{JAVA_OPTS}"

export CLASSPATH="<$CATALINA_HOME_or_$CATALINA_BASE>/compress/ebx-
lz4.jar:$CLASSPATH"
```

Where `<$CATALINA_HOME_or_$CATALINA_BASE>` must be replaced by `$CATALINA_HOME` or `$CATALINA_BASE` if they have been configured. Otherwise this piece of text must be replaced by the Tomcat installation directory's path.

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`.

Note:

For a [Data compression library](#) native installation, ensure to only reference it in the *CLASSPATH* environment variable.

EBX® and third-party libraries deployment

1. Copy third-party libraries from the *EBX® CD* to `$CATALINA_HOME/lib/` (or `$CATALINA_BASE/lib/`) directory, except for the [Data compression library](#). 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/h2-2.1.212.jar` (default persistence factory) coming from `ebx.software/lib/lib-h2` directory.

The exact description of these components is given in chapter [Software components](#).

Obviously, if those components are already deployed on the class-loading system, they do not have to be duplicated.

2. Create a directory dedicated to the [Data compression library](#) (for example `$CATALINA_HOME/compress` or `$CATALINA_BASE/compress`) and copy it there.

Note:

Ensure that the library is copied in the directory pointed out by the previously updated *CLASSPATH* environment variable.

3. 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`

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

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](#) helps to configure the default properties of the initial repository.

Installation note for WebLogic 14c

- [Overview](#)
- [Requirements](#)
- [WebLogic Application Server installation](#)
- [EBX® home directory configuration](#)
- [WebLogic Application Server and Java Virtual Machine configuration](#)
- [EBX® and third-party libraries deployment](#)
- [Data source and JDBC provider configuration](#)
- [EBX® application deployment](#)
- [EBX® application start](#)

Overview

Attention:

- 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](#).
- To avoid unpredictable behavior, the guideline to follow is to avoid any duplicates of `ebx.jar`, `ebx-lz4.jar` or other libraries in the class-loading system.

Requirements

- Certified Oracle Java SE 11 ([update 6 and above](#))
- WebLogic Server 14c
- Database and JDBC driver
- EBX® CD

See also:

- [Supported environments](#)

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:
 - a. Create Domain: choose 'Create a new domain' and specify the domain home directory, then click 'Next'
 - b. Templates: keep as default and click 'Next'

- c. Administrator Account: enter a domain administrator username and password and click 'Next'
- d. Domain Mode and JDK: choose the production mode and your JDK installation home and click 'Next'
- e. Advanced configuration: check 'Administration server' and 'Topology'. That way, we create two independent domain nodes: an administration one and an application one.

Click 'Next'

- f. Administration Server: enter your administration node name (for example 'AdminServer') and listen port (by default `7001`), then click 'Next'
- g. Managed Servers: add the application node name (for example 'EbxServer') and listen port (for example `7003`), then click 'Next'
- h. Clusters: keep as default and click 'Next'
- i. Server Templates: keep as default and click 'Next'
- j. Machines: keep as default and click 'Next'
- k. Configuration Summary: click 'Create'
- l. Configuration Process: click 'Next'
- m. End Of Configuration: click 'Finish'

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.

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}"
```

2. Edit the `<DOMAIN_HOME>/bin/setDomainEnv.sh` script file by adding the following line:

```
PRE_CLASSPATH="<path_to_the_data_compression_library>"

# For our example

# PRE_CLASSPATH="${DOMAIN_HOME}/compress/ebx-lz4.jar"
```

EBX® and third-party libraries deployment

1. Copy third-party libraries from the *EBX® CD* to the `<DOMAIN_HOME>/lib` directory except for the [Data compression library](#). In our example, for an H2 standalone data base, we will have:

`<DOMAIN_HOME>/lib/h2-2.1.212.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](#). 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).

2. Create a directory dedicated to the [Data compression library](#) (for example `<DOMAIN_HOME>/compress`) and copy it there.

Note:

Ensure that the library is copied in the directory pointed out by the previously updated `PRE\_CLASSPATH` environment variable.

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':

- a. Set: Type Name: `EBX_REPOSITORY`, JNDI Name: `EBX_REPOSITORY`, Database Type: *Your database type*

Click 'Next'

- b. Choose your database driver type, and click 'Next'
- c. Uncheck 'Supports Global Transactions', and click 'Next'
- d. Setup your database 'Connection Properties' and click 'Next'
- e. Click 'Test Configuration' and then 'Finish'
- f. Switch on the 'Targets' tab and select all Servers, then click 'Save'
- g. Restart the Administration server (for example 'AdminServer'), using:

```
<DOMAIN_HOME>/bin/stopWebLogic.sh
```

```
<DOMAIN_HOME>/bin/startWebLogic.sh
```

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':
 - a. 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'
 - b. Choose the installation type and scope: Click on 'Install this deployment as an application', 'Global' default scope and click 'Next'
 - c. Select the deployment targets: Select a node (for example 'EbxServer') from the 'Servers' list and click 'Next'
 - d. 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
```

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](#) helps to configure the default properties of the initial repository.

TIBCO EBX® main configuration file

- [Overview](#)
 - [Location of the file](#)
 - [Custom properties and variable substitution](#)
- [Setting automatic installation on first launch](#)
- [Setting the EBX® root directory](#)
- [Configuring the EBX® repository](#)
- [Configuring the user and roles directory](#)
- [Configuring EBX® localization](#)

- Setting temporary files directories
- Activating the audit trail
- Activating the legacy XML audit trail (deprecated)
- Configuring the EBX® logs
 - Custom 'ebxFile' appender
 - Custom SMTP appender
 - Custom module log threshold
 - Add-on module log threshold
- Activating and configuring SMTP and emails
- Configuring data services
- Activating and configuring JMS
- Configuring distributed data delivery (D3)
- Configuring REST toolkit services
- Activating and configuring staging
- Configuring Web access from end-user browsers
 - HTTP Authorization header policy
 - URLs computing
 - Proxy mode
 - Reverse-proxy mode
- Configuring failover
- Tuning the EBX® repository
- Miscellaneous
 - Activating data workflows
 - Disabling user task legacy mode
 - Disabling hierarchy plan view
 - Activating delegation to launch a workflow by API
 - Log procedure starts
 - Log validation starts
 - Request cache activation

- Request duration threshold for logs
- Request repetition threshold for logs
- Request optimization threshold for logs
- Request SQL representation for logs
- Request optimization for sorting a table by foreign key labels
- Deployment site identification
- Dynamically reloading the main configuration
- Declaring modules as undeployed
- Module public path prefix
- EBX® run mode
- Resource filtering
- Validation report page
- Validation report logs

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](#). Administrators can also perform further configuration through the user interface, which is then stored in the EBX® repository.

See also:

- [Deployment details](#)
- [UI administration](#)

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.

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).
```

```
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 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 non-encrypted password used for the 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
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 <login>
<password_to_encrypt>

ebx.install.admin.password.encrypted=ff297ae08f7eeb63230b55f7c45a720a017bc
71d22eaaec...
```

Setting the EBX® root directory

The EBX® root directory contains the Lucene indexes directory, the archives and, when the repository is persisted on H2 standalone mode, the H2 database files. It also contains the legacy XML audit trail (deprecated).

```
#####

## Path for EBX® XML repository

#####

ebx.repository.directory=${ebx.home}/ebxRepository
```

See also:

- [Monitoring and clean up of the file system](#)

Configuring the EBX® repository

Before configuring the persistence properties of the EBX® repository, carefully read the section [Technical architecture](#) in the chapter 'Repository administration'.

The required library (driver) for each supported database is described in the chapter [Database drivers](#).

See also:

- [Repository administration](#)
- [Rules for the database access and user privileges](#)
- [Supported databases](#)
- [Data source of the EBX® repository](#)
- [Database drivers](#)

```
#####

## 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 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=xxxxxxxxx

#ebx.persistence.password=yyyyyyyyy

#####

## Case EBX® persistence system is Oracle database.

#####

```

```

#ebx.persistence.factory=oracle

## Specific properties to be set 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=xxxxxxxxx

#ebx.persistence.password=yyyyyyyyy

## 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 SAP Hana

#####

#ebx.persistence.factory=hana

## Specific properties to be set 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:sap://127.0.0.1:39041

#ebx.persistence.driver=com.sap.db.jdbc.Driver

#ebx.persistence.user=xxxxxxxxx

#ebx.persistence.password=yyyyyyyyy

#####

## Case EBX® persistence system is Microsoft SQL Server.

#####

#ebx.persistence.factory=sqlserver

## Specific properties to be set 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;databasename=ebxDatabase

#ebx.persistence.driver=com.microsoft.sqlserver.jdbc.SQLServerDriver

#ebx.persistence.user=xxxxxxxxx

#ebx.persistence.password=yyyyyyyyy

#####

```



```

## Case EBX® persistence system is Microsoft Azure SQL database.

#####

#ebx.persistence.factory=azure.sql


## Specific properties to be set 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=xxxxxxxxx

#ebx.persistence.password=yyyyyyyyy

#####

## Case EBX® persistence system is PostgreSQL.

#####

#ebx.persistence.factory=postgresql


## Specific properties to be set 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=xxxxxxxxx

#ebx.persistence.password=yyyyyyyyy
```

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](#)
- `DirectoryFactory`

```
#####

## 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
```

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](#)

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 = /tmp/java

ebx.temp.directory = ${java.io.tmpdir}

## 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
```

Activating the audit trail

By default, the audit trail logging is deactivated. It can be activated by configuring the "audit" log category to log INFO level messages. For instance the following will log audit trail messages using a specific `'ebxFile'` appender:

```
ebx.log4j.category.log.audit = INFO, ebxFile:audit
```

See also:

- [Configuring the EBX® logs](#)
- [Audit trail](#)

Activating the legacy XML audit trail (deprecated)

By default, the XML audit trail is deactivated. It can be activated using the following variable:

```
#####

## The XML history has been replaced by an SQL history.

## This old XML history can be activated using the following variable.

## Default is false.

#####

ebx.history.xmlaudittrail.activated = false
```

See also:

- [Legacy XML Audit trail](#)

Configuring the EBX® logs

The most important logging categories are:

```
ebx.log4j.category.log.kernel
```

Logs for EBX® main features, processes, exceptions and compilation results of modules and data models.

ebx.log4j.category.log.workflow	Logs for main features, warnings and exceptions about workflow.
ebx.log4j.category.log.persistence	Logs related to communication with the underlying database.
ebx.log4j.category.log.setup	Logs for the compilation results of all EBX® objects, except for modules and data models.
ebx.log4j.category.log.validation	Logs for datasets validation results.
ebx.log4j.category.log.mail	Logs for the activity related to the emails sent by the server (see Activating and configuring SMTP and emails). Note: This category must not use the Custom SMTP appender in order to prevent infinite loops.
ebx.log4j.category.log.d3	Logs for D3 events on EBX®.
ebx.log4j.category.log.dataservices	Logs for data service events in EBX®.
ebx.log4j.category.log.monitoring	Raw logs for memory monitoring .
ebx.log4j.category.log.request	Logs for <code>Request</code> and <code>Query</code> events in EBX®: executions having a duration exceeding <code>ebx.logs.request.durationThreshold</code> milliseconds, requests / queries where the optimization phase exceeds <code>ebx.logs.request.optimizationThreshold</code> milliseconds and executions that internally throw an exception while iterating the results. All requests / queries are logged regardless of their duration, if log level is set to <code>DEBUG</code>.

ebx.log4j.category.log.restServices	Logs for REST services events in EBX®, including those from the REST Toolkit .
ebx.log4j.category.log.audit	Logs for the audit trail feature.

Some of these categories can also be written to through custom code using the `LoggingCategory` 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
out

#-----

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

ebx.log4j.category.log.audit=INFO, ebxFile:audit

```

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`, with a threshold set to `DEBUG`.

The property `ebx.log4j.appender.ebxFile.backup.Threshold` allows defining the maximum size (in megabytes) 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 size (in megabytes) 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](#)

Custom module log threshold

By default, the log level threshold of the logging category associated with a custom module is set to `INFO`.

This threshold can be customized by setting the property

`ebx.log4j.category.log.wbp.xxxxxx` for the custom module `xxxxxx`.

Example: `ebx.log4j.category.log.wbp.mycompany-module=DEBUG`.

See also:

- `ModuleContextOnRepositoryStartup.getLoggingCategory`

Add-on module log threshold

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` for the add-on module `ebx-addon-xxxxxx`.

Example: `ebx.log4j.category.log.addon.daqa=DEBUG`

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](#)

```
#####

## 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 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=

## SMTP socket connection timeout value in milliseconds (default is
600000).

#ebx.mail.smtp.connectionTimeout=600000

## SMTP socket read timeout value in milliseconds (default is 600000).

#ebx.mail.smtp.timeout=600000

## SMTP socket write timeout value in milliseconds (default is 600000).

#ebx.mail.smtp.writeTimeout=600000


## Activate SSL (true or false, default is false).

## If SSL is activated, a SSL factory and a SSL provider are required.

#ebx.mail.smtp.ssl.activate=true
```

```
#ebx.mail.smtp.ssl.provider=com.sun.net.ssl.internal.ssl.Provider

#ebx.mail.smtp.ssl.factory=javax.net.ssl.SSLSocketFactory
```

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

## Specifies the default maximum pagination size value for the select
```

```
## operations. This configuration is used by SOAP and REST connectors.
```

```
## Default value is 10000, maximum recommended value is 100000
```

```
#ebx.dataservices.pagination.maxSize.default= 10000
```

```
## Specifies the default pagination size value for the select
```

```
## operations. This configuration is used by the SOAP connector.
```

```
## Default value is 10.
```

```
#ebx.dataservices.pagination.pageSize.default=10
```

```
## 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, when a REST data service authentication negotiate fails,

## EBX response includes fallback to 'Basic' authentication method by
setting

## the HTTP header 'WWW-Authenticate' to 'Basic'.

## Note: This property only activate/deactivate

## the authentication fallback.

## 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

```

Activating and configuring JMS

See also:

- [JMS for data services](#)

```
#####

## 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
```

Configuring distributed data delivery (D3)

See [Configuring D3 nodes](#) for the main configuration file properties pertaining to D3.

See also:

- [JMS for distributed data delivery \(D3\)](#)
- [Introduction to D3](#)

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
```

Activating and configuring staging

Staging is enabled/disabled via configuration. Some options can be set so as to optimize memory and disk usage. The properties are configured as follows:

```
#####

## Staging configuration

#####

## Activates staging (true or false, default is true).

#ebx.staging.activated=true

## Defines the max size on disc

## dedicated to staging temporary folders.

## 0 defines an infinite size.

## If unset, the default value is 0.

#ebx.staging.maxTemporaryFolderSizeInBytes=0

## Max size for canonicalization of attachments in bytes.

## Canonicalization consumes memory and CPU on large xml and json for
attachments.

## If not set, default value is 10 Mo (10485760).

#ebx.staging.maxCanonicalizationSizeInBytes=10485760

## Defines the maximum number of staging element per list when browsing a
```

```
repository
```

```
## If not defined, the default value is 1000.
```

```
#ebx.staging.maxElementPerList=1000
```

See also:

- [Staging](#)

Configuring Web access from end-user browsers

HTTP Authorization header policy

EBX® natively offers three policies to send and receive credentials using HTTP headers:

<code>standard</code>	It corresponds to the authentication scheme, using the HTTP Authorization header, described in the RFC 2617 .
<code>ebx</code>	To prevent HTTP Authorization header override issues, this policy acts the same as the <code>standard</code> but the credentials are stored in an EBX® specific HTTP header.
<code>both</code>	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 \(deprecated\)](#)

```

#####

## 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

##

## 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`.

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`.

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](#)

```
#####
```

```
## 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

```

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
```

See also:

- [Validation report page](#)

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 true.

#####

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 `false`.

See `UserTask.UserTaskMode.LEGACY_MODE` for more information.

```
#####

## Disables legacy work item mode (default is false)

## Specify if the creation service of user task in legacy mode must be
offered

## in workflow modeling.

#####

#ebx.manager.workflow.legacy.userTaskMode=true
```

Disabling hierarchy plan view

This parameter specifies whether the hierarchy plan view is hidden. The default value is `true`.

```
#####

## Activate or deactivate Workflow hierarchy plan view

#####

ebx.manager.workflow.hierarchyPlanView.hidden=false
```

Activating delegation to launch a workflow by API

This parameter activates delegation to launch a workflow by API. The default value is `false` which means that permissions will be always checked unless this parameter is activated.

```
#####

## Activate delegation to launch workflows by API

#####

ebx.workflow.api.processLauncher.permissionCheckDelegated=true
```

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
```

Request cache activation

This parameter specifies whether or not to activate the cache for the internal request-to-query conversions that are required by the query engine when a **Request** is executed. The activation of this cache will improve the performance on certain types of requests, especially if they are executed multiple times.

```
# Activates the request cache, which can improve the performance of
certain requests. Possible values are DISABLED,

# WITH_CHECK (ensures that a cached request executes the same query plan
as its non-cached version, recommended),

# EXCEPTION (throws an exception if the query plan for a request differs
from its non-cached version), and ENABLED.

# The default value is WITH_CHECK.
```

```
#ebx.cache.request.mode=WITH_CHECK
```

Request duration threshold for logs

This parameter specifies in milliseconds the threshold of duration of **Request** and **Query** to be logged. Logs are generated if logging category `ebx.log4j.category.log.request` level is not higher than **INFO**. If the level is **DEBUG**, all **Request** and **Query** are logged.

```
#####

## Specifies in milliseconds the threshold of duration of Requests and
## Queries

## to be logged

## Default value is 1000 ms.

## If unset, the default value is used.

#####

#ebx.logs.request.durationThreshold=1000
```

Request repetition threshold for logs

This parameter specifies in milliseconds the delay between 2 logs for **Request** and **Query** that goes beyond the threshold of duration. If this value is greater than 0, and the query duration goes beyond the threshold of duration, it will be logged again repeatedly with at least this delay between each log. As log messages include duration, this is useful to track long queries duration.

```
#####

## Specifies in milliseconds the delay between two log entries for
## Requests and

## Queries that goes beyond the threshold of duration. If this value is
## greater

## than 0, and the query duration goes beyond the threshold of duration,
## it will
```



```
## be logged again repeatedly with at least this delay between each log.

## Default value is 30000 ms.

#####

#ebx.logs.request.logAgainEvery=30000
```

Request optimization threshold for logs

This parameter specifies in milliseconds the threshold of the optimization phase of `Request` and `Query` to be logged. Logs are generated if logging category `ebx.log4j.category.log.request` level is not higher than `INFO`.

```
#####

## Specifies in milliseconds the threshold of optimization of Requests and
## Queries

## to be logged. Default value is 1000 ms.

#####

#ebx.logs.request.optimizationThreshold=1000
```

Request SQL representation for logs

This parameter specifies if the logs for a certain `Request` shall include the corresponding SQL representation of the request. Notice that this conversion takes a toll in terms of performance, and not all requests can be converted into SQL. The default value is `false`.

```
#####

## Specifies if the logs of requests shall include the SQL representation.

## Default is false.

#####
```

```
#ebx.log4j.category.log.request.enableSQLConversion=false
```

Request optimization for sorting a table by foreign key labels

This parameter specifies whether the query optimizer rewrites certain combinations of sort and join operations to a specialized operator that carries out the join(s) in a sorted manner.

```
#####

## Properties to activate/deactivate EnumerableSortedNestedLoopLeft/Right
operators,

## which can improve the performance of requests sorted by foreign key
labels.

## Default value is true.

#####

#ebx.query.sortedNestedLoopLeft.enabled=true

#ebx.query.sortedNestedLoopRight.enabled=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 undeployed

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](#)
- [Dynamically reloading the main configuration](#)

```
#####

## 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](#) 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](#) 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):

Documentation pane	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.
Compilation information	Module and schema compilation information is displayed in the dataset validation report.
Java bindings	The generation of Java bindings is available if the schema of the dataset mentions at least one binding.
Web component link generator	The Web component link generator is available on datasets and dataspace.
Data model assistant	Data model configuration and additional options, such as Services, Business Objects and Rules, Java Bindings, Toolbars and some advanced properties.
Workflow modeling	Declare specific script tasks.

Log	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.
Product documentation	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 can be of type [pattern] or "m:[pattern]:".
```



```
## the list can be void

#####

ebx.resource.exclude=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=200
```

See also:

- [Tuning the EBX® repository](#)

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
```

By default, the content of the validation of a dataset or a table is logged. Logging the content of the validation reports can be deactivated using the following property:

```
#####

## Specifies whether the content of validation reports is logged when
validating

## a dataset or a table.

## Default is true.

#####

ebx.validation.report.logContent=true
```

See also:

- [Tuning the EBX® repository](#)

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](#) 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](#).

Note:

The EBX® main configuration file must also be properly configured. See [TIBCO EBX® main configuration file](#).

- Configuration steps
 - Validating the license agreement
 - Configuring the repository
 - Defining users in the default directory
 - Validating the information entered
 - Installing the EBX® repository

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.

Id of the repository (`repositoryId`)

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.
Repository label	Defines a user-friendly label that indicates the purpose and context of the repository.

See also:

- [TIBCO EBX® main configuration file](#)

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](#)

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.

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.

- [Deploying an add-on module](#)
- [Registering an add-on module](#)
- [Activating an add-on module](#)
- [Deleting an add-on module](#)

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.

Copy the add-on common JAR file (named `lib/ebx-addons.jar`) in the EBX® class-path; it must be accessible from the ebx.jar class-loader.

Copy the add-on common WAR file (named `war/ebx-addons-common.war`) in the same folder or EAR as EBX® built-in web applications.

If an add-on is used, copy the EBX® add-on WAR file (named `war/ebx-addons-<name>.war`) in the same folder or EAR as EBX® built-in web applications.

See [Web applications](#), [Deployment details](#) and [Installation notes](#) for more information.

Note:

The add-on log level can be managed in the [main configuration file](#).

Registering an add-on module

Registering an add-on makes its configuration available in the admin section. Add-on features are only available to end-users when the add-on is also [activated](#).

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.
5. Click on **Save**.

Note:

Unregistering an add-on will not delete any existing configuration, but will make it available in the UI until the add-on is registered again.

Activating an add-on module

Activating an add-on makes its features available to the end-users. Only registered add-ons can be activated.

To activate an 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. Select the registered add-on you are activating and enable the 'Activation' field.
4. Click on **Save**.

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.

EBX® Container Edition

- [Building the image](#)
- [Running the image](#)
- [Customizing the image](#)

Building the image

- [Overview](#)
- [Requirement](#)
- [Building the image](#)
 - [Download installers](#)
 - [Running the interactive installer](#)
 - [Running the batch installer](#)
 - [Testing the image](#)
 - [Sharing the image](#)
 - [Other architectures than Intel or AMD](#)

Overview

TIBCO EBX® Container Edition image is Linux based (**amd64** architecture) and includes the following:

- Application Server Apache Tomcat® 9.0

- Java JDK 17

Requirement

The image can be built on a POSIX native system or using a compatible layer:

- **Operating system** : Linux, macOS, Windows (10 version 2004, Server 2022 or greater) using WSL2,
- **Processor type** : Intel or AMD 64 bits,
- **Docker Engine** installed and running,
- Access to **Internet** .

Building the image

Download installers

To build an EBX® image, one needs to download file

TIB\ebx\6.1.0\addon\6.1.X\container\edition.amd64.zip from TIBCO eDelivery.

Running the interactive installer

To start installer on **Linux** or **macOS** :

- Unzip the **TIB\ebx\6.1.0\addon\6.1.X\container\edition.amd64.zip** ,
- Open a terminal in folder where file **ebx-ce-installer.sh** file is located,
- Execute command **./ebx-ce-installer.sh** .

To start installer on **Windows** :

- Unzip the **TIB\ebx\6.1.0\addon\6.1.X\container\edition.amd64.zip** ,
- Open a Windows PowerShell in folder where file **ebx-ce-installer.bat** file is located,
- Execute command **ebx-ce-installer.bat** .

Follow instructions and select **addons** to be added to the image.

The installer will then build the image and print a summary similar to:

```
*****
*****
```

The TIBCO EBX Container Edition image was successfully created with

following names:

ebx:latest

ebx:6.1.0-dama-dint-dmdv-dpra-mame-tese-6.1.X

This image includes the following addon(s):

Digital Asset Manager (DAMA)

Data Exchange (New) (DINT)

Data Model and Data Visualization (DMDV)

Insight (New) (DPRA)

Match and Merge (MAME)

Information Search (TESE)

To run this image with the default configuration and an embedded database, use command:

```
docker run -p 8080:8080 -d ebx:latest
```

You can test EBX by visiting <http://localhost:8080> in a browser.

To run this image with other configurations, for example with an external database, see

documentation.

```
*****
*****
```

Running the batch installer

To start batch installer on **Linux** or **macOS** :

- Unzip the **TIB\ebx\6.1.0\addon\6.1.X\container\edition.amd64.zip** ,
- Open a terminal in folder where file **ebx-ce-installer-batch.sh** file is located,
- Execute command **./ebx-ce-installer-batch.sh**

Batch script to build image with optionally addon(s).

```
-a                add all addons

-t <tag_file>    create a file named tag_file including the image tag
built

--help          display this help and exit

[addon]...      add specific addons (separate by a space)

                for sample: mame tese
```

To start batch installer on **Windows** :

- Unzip the **TIB\ebx\6.1.0\addon\6.1.X\container\edition.amd64.zip** ,
- Open a Windows PowerShell in folder where file **ebx-ce-installer-batch.bat** file is located,
- Execute command **ebx-ce-installer-batch.bat ...**

The batch installer will then build the image and print a summary identical to the interactive installer.

Testing the image

The image can be run locally using command:

```
docker run -p 8080:8080 -d ebx:latest
```

In **production** , it is recommended to **not** use tag **latest** . The installer will always generate another tag depending on selected addons.

If previous sample, the tag is **6.1.0-dama-dint-dmdv-dpra-mame-tese-6.1.X** because the following addons where selected:

- Digital Asset Manager (DAMA),
- Data Exchange (New) (DINT),
- Data Model and Data Visualization (DMDV),

- Insight (New) (DPRA),
- Match and Merge (MAME),
- Information Search (TESE).

The image can then be run using command:

```
docker run -p 8080:8080 -d ebx:6.1.0-dama-dint-dmdv-dpra-mame-tese-6.1.X
```

If no addons are selected, the tag is **6.1.0**

The image can be run using command:

```
docker run -p 8080:8080 -d ebx:6.1.0
```

Sharing the image

The steps to share an image depends on customer's company infrastructure.

In following example, the image is pushed to a Docker private registry named **myregistry** :

```
docker tag ebx:6.1.0 myregistry:5000/ebx:6.1.0

docker push myregistry:5000/ebx:6.1.0
```

Other architectures than Intel or AMD

The installers are tested on Intel or AMD 64 bits systems. On these platforms the generated image's architecture will be **linux/amd64** .

The installers may succeed when running on a system using another processor. In that case the generated image architecture will be linux for that processor type.

For example, if the installer runs on a MacOS M1 workstation, the generated image's architecture will be **linux/arm64** .

Attention

Generating an image on a system that is not Intel or AMD is **experimental** and is **not supported in production** .

Running the image

- Support Policy
 - Docker
 - Kubernetes
- Starting EBX®
 - First-launch assistant
 - Automatic initialization
 - Supported browsers
- Container access
- Environment variables
 - Disabling First-launch assistant
 - Automatic repository installation on first launch
 - URL configuration
 - Database connectivity
 - Email connectivity
 - Memory configuration
 - Authentication for REST services
 - Staging
- Configuration files
 - File ebx-default.properties
 - File ebx-container.properties
- Volumes
- Linux user and group
- Host configuration
- Logs access

Support Policy

Docker

TIBCO EBX Container Edition is tested with the Docker Engine version 20.10.

Kubernetes

TIBCO EBX Container Edition is tested with one of the latest stable Kubernetes versions available at the time of product testing before the release.

Current TIBCO EBX Container Edition release was tested with a certified Kubernetes implementation version 1.23.

Starting EBX®

First-launch assistant

To start EBX® with default configuration that includes an embedded H2 database, execute command:

```
docker run -p 8080:8080 -d ebx:6.1.0
```

Using a browser, you can connect to EBX® with URL <http://localhost:8080> . This will display the **first-launch assistant** that will help you configure EBX®.

For more information on the first launch assistant, see chapter [Initialization and first-launch assistant](#) .

Automatic initialization

To start EBX® with automatic initialisation on first startup and an embedded H2 database, execute command:

```
docker run -d -p 8080:8080 \  
  
-e "EBX_FLA_DISABLED=true" \  
  
-e "EBX_INSTALL_ADMIN_PASSWORD=<password>" \  
  
ebx:6.1.0
```

The EBX® repository will be automatically created on first startup.

Using a browser, you can connect to EBX® with URL <http://localhost:8080> . This will display the EBX® login screen. The username for administrator is **admin** and the password is the one specified in previous command.

Note

It's possible to specify another username for the administrator. For more information see [Automatic repository installation on first launch](#) .

Supported browsers

For details a supported browsers see: [Supported Web Browsers](#) .

Container access

The following command will start a bash shell inside the EBX® container:

```
docker exec -it <container-id> bash
```

Environment variables

This chapter describes the environment variables supported by **EBX® Container Edition** .

All are optional.

Disabling First-launch assistant

For security reasons, one might want to disable the first-launch assistant in all circumstances.

This is achieved by setting environment variable to **EBX\FLA\DISABLED** to **true** .

Automatic repository installation on first launch

If the repository is not yet initialized and first-launch assistant is disabled, EBX® will automatically trigger its installation if following mandatory variables are provided:

NAME	DEFAULT	DESCRIPTION
EBX\INSTALL\ADMIN\LOGIN	admin	Sets the EBX® administrator login name. This parameter is ignored if repository variable EBX\FLA\DISABLED

NAME	DEFAULT	DESCRIPTION
		value is not true or if repository is already initialized.
EBX_INSTALL_ADMIN_PASSWORD		Sets the EBX® administrator login name. This parameter is mandatory if EBX_FLA_DISABLED value is true and is ignored if repository variable EBX_FLA_DISABLED value is not true or if repository is already initialized.

Note

If mandatory variables are not provided, EBX® will display an error message.

Example

To automatically install repository launch EBX® using following command:

```
docker run -d -p 8080:8080 \
  -e "EBX_FLA_DISABLED=true" \
  -e "EBX_INSTALL_ADMIN_LOGIN=<login-name>" \
  -e "EBX_INSTALL_ADMIN_PASSWORD=<password>" \
  ebx:6.1.0
```

URL configuration

Some EBX® features require generating URLs. Specific configuration may be required to achieve this, for example if EBX® is running behind a reverse proxy or on a Kubernetes cluster.

NAME	DEFAULT	DESCRIPTION
EBX_IS_SECURED	If incoming request is HTTPS, "true" is assumed or else "false" is assumed.	If "true", the protocol "HTTPS" is assumed. If "false", the protocol "HTTP" is assumed.
EBX_HOSTNAME	The host name specified by the incoming HTTP(S) request.	The EBX® server host name.
EBX_PORT	The port number specified by the incoming HTTP(S) request.	The EBX® server port number.
EBX_ROOT_PATH	By default, the context path is empty.	If set, all EBX® urls will be prefixed by this value. The value must have a leading / and must not have a trailing / except if value is / . For example a valid value is / mdm/sales . Setting this variable is useful when running more than one instance of EBX with the same host name.
EBX_URL_DEFAULT		This environment variable is used when a background task needs to calculate a URL to EBX. It should be set to a full URL without the path component (EBX_ROOT_PATH applies for the path component). For example a valid value is: https://host_name . If EBX_URL_DEFAULT is not specified and EBX_HOSTNAME is specified, a default is calculated with following assumptions: - If

NAME	DEFAULT	DESCRIPTION
		EBX_IS_SECURED is not specified, 'false' is assumed, - If EBX_PORT is not specified, 443 is assumed if EBX_IS_SECURED is true or else 80 is assumed.

Database connectivity

For information on supported databases see chapter [Supported databases](#) .

By default, an embedded H2 database is used. Data for this H2 database is persisted at location **/ebx/data/h2** .

An external database may be configured using following variables:

NAME	TIBCO EBX® MAIN CONFIGURATION FILE EQUIVALENT	DESCRIPTION
EBX_DB_FACTORY	ebx.persistence.factory	Specifies the type of database server.
EBX_DB_URL	ebx.persistence.url	The JDBC URL. It has format: jdbc:postgresql://<database_host>:<database_port_number>:<database_name> .
EBX_DB_USER	ebx.persistence.user	The database user id.
EBX_DB_PASSWORD	ebx.persistence.password	The database user password.

For more information on these variables see their TIBCO EBX® main configuration file equivalent in chapter [Configuring the EBX® repository](#) .

Note

The container includes JDBC drivers only for H2, PostgreSQL and Microsoft SQL Server. Using other databases that are supported by EBX® require adding the driver.

For instructions on how to add a driver, see [Adding a new JDBC driver](#) .

Example

To start EBX® that connects to a Postgres database, execute following command:

```
docker run -d -p 8080:8080 \

-e "EBX_DB_FACTORY=postgresql" \

-e "EBX_DB_URL=jdbc:postgresql://<server_name>:5432/<database_name>" \

-e "EBX_DB_USER=<user_name>" \

-e "EBX_DB_PASSWORD=<user_name>" \

ebx:6.1.0
```

Email connectivity

The EBX Mail service can be configured through the following environment variables :

NAME	TIBCO EBX® MAIN CONFIGURATION FILE EQUIVALENT	DEFAULT	DESCRIPTION
EBX_SMTP_HOST	ebx.mail.smtp.host		SMTP server host name.
EBX_SMTP_PORT	ebx.mail.smtp.port		SMTP server port number.
EBX_SMTP_CONNECTION_TIMEOUT	ebx.mail.smtp.connectionTimeout	60000	SMTP socket connection timeout value in milliseconds.
EBX_SMTP_TIMEOUT	ebx.mail.smtp.timeout	600000	SMTP socket read timeout value in milliseconds.
EBX_SMTP_WRITE_TIMEOUT	ebx.mail.smtp.writeTimeout	600000	SMTP socket write timeout value in milliseconds.

NAME	TIBCO EBX® MAIN CONFIGURATION FILE EQUIVALENT	DEFAULT	DESCRIPTION
EBX_SMTP_LOGIN	ebx.mail.smtp.login		SMTP server login id.
EBX_SMTP_PASSWORD	ebx.mail.smtp.password		SMTP server login password.
EBX_SMTP_SSL_ENABLED	ebx.mail.smtp.ssl.activate	true	Enables SSL. Value can be 'true' or 'false'.
EBX_WORKFLOW_MAIL_SENDER	ebx.workflow.mail.sender		The workflow sender email. If not set, Workflows cannot send notifications.

More information on the used properties can be found in chapter [Activating and configuring SMTP and emails](#) .

Memory configuration

Environment variables **JAVA_MEMORY_PERCENT** may be used to configure the percentage of the container memory that is assigned to the JVM the runs EBX®. It must be an integer value between 0 and 100.

If not set, a default value is used at startup.

Note

This variable is for advanced usage. Setting it too low or too high may cause runtime issues.

Authentication for REST services

Basic authentication for REST services is not enabled by default.

To enable this feature, set environment variable **EBX_REST_AUTHENTICATION_BASIC** to **true** .

Staging

The environment variable **EBX\STAGING\ACTIVATED** activates, if value is **true** , or deactivates, if value is **false** , the staging feature.

By default, staging is activated.

Configuration files

Two Java property files are currently used to configure EBX®.

On startup EBX® reads property files in the following order:

- **/opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties**
- **/my_custom/conf/ebx-container.properties**

File ebx-default.properties

The file **/opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties** sets default EBX® configuration properties for the container.

It should **never** be modified at runtime as this may prevent easily updating EBX® to a next version, instead use **/opt/ebx/conf/ebx-container.properties** .

File ebx-container.properties

The file **/opt/ebx/conf/ebx-container.properties** is by default empty. Any property value specified here will override the value set by **ebx-default.properties** .

This file is useful to change a property at runtime. To change a property at run time, create a new file, for example **/my_custom/conf/ebx-container.properties** , containing the new property values and mount de parent folder from the host to the container:

```
docker run -v /my-custom/conf:/opt/ebx/conf -p 8080:8080 -d ebx:6.1.0
```

For the list of properties supported by EBX® see chapter [TIBCO EBX® main configuration file](#) .

Volumes

This image defines the following volumes:

LOCATION	DESCRIPTION
/ebx/data	The EBX® root directory is located in this volume. It contains EBX® indexes and, when H2 embedded database is used, persisted data. It is recommended to use a high performance volume, for example an SSD. Using a shared network volume, for example a NFS mount, is not recommended and may cause performance and stability issues.
/ebx/logs	This volume is used for log files.
/ebx/temp	This volume is used for temporary files.

Note

The volume **/ebx/data** should be mapped to a persistent volume even when an external database is used. If not, EBX will have to rebuild its indexes on startup which may considerably increase boot time.

Linux user and group

The Linux process running EBX is owned by user **ebx** (uid 1500).

User ebx's primary group is **ebx** (guid 1500).

Host configuration

It may be necessary to configure the Host so that the EBX Container can reserve the resources required by numerous memory-mapped files.

On a Linux OS, this can be done by issuing the following commands:

```
ulimit -n 512000

sysctl vm.max_map_count=262144
```

If more than one EBX container may run on the same host at the same time, one needs to increase these values accordingly.

On a Kubernetes cluster, setting these values can be done by running an init container in a privilege security context. The following sample is a partial Kubernetes deployment YAML file that sets **vm.max_map_count** :

```
apiVersion: apps/v1

kind: Deployment

...

spec:

  ...

  template:

    ...

    spec:

      initContainers:

        - name: sysctl

          image: busybox

          securityContext:

            privileged: true

            runAsUser: 0

          command: ['sh', '-c', 'sysctl -w vm.max_map_count=262144']

      containers:

        - name: "ebx-pod"

          image: "ebx"

          ...
```

Logs access

Logs are sent to the stdout and stderr output streams and can be viewed using the following command:

```
docker logs <container-id>
```

Logs for both EBX® and Tomcat will be displayed.

Log files are also available under folder **/ebx/logs** :

- EBX® logs files are in folder **/ebx/logs/ebx**
- Tomcat logs files are in folder **/ebx/logs/tomcat** .

Customizing the image

The EBX® Container Edition image can be used as a parent image to create a customized image.

- [Setting default configuration](#)
- [Adding a custom module](#)
- [Adding a new locale](#)
- [Adding a new JDBC driver](#)

Setting default configuration

Setting default EBX® configuration should not be based on **/opt/ebx/conf/ebx-container.properties** as this file may be overridden at runtime.

Instead, proceed as following in the Docker file:

- Rename file **/opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties** , for example to **ebx-default-original.properties** .
- Create a new file **/opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties** container new property values. This file must define property **ebx.file.previous** set to the original property file new name, for example **ebx-default-original.properties** .

For the list of properties supported by EBX® see chapter [TIBCO EBX® main configuration file](#) .

Here is a sample Docker file that set the locale to "en-US"

```
FROM ebx:6.1.0

RUN mv /opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties \

/opt/ebx/webapps/ebx/WEB-INF/ebx-default-original.properties

RUN echo "ebx.file.previous=ebx-default-original.properties" >> \

/opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties

RUN echo "ebx.locales.available=en-US" >> /opt/ebx/webapps/ebx/WEB-INF/
ebx-default.properties
```

Adding a custom module

One can extend EBX® by developing custom modules. An EBX® module is a standard Java EE web application, packaging various resources such as XML Schema documents, Java classes and static resources.

For more information on EBX® module see chapter [Packaging TIBCO EBX® modules](#) .

With EBX® Container Edition, it is recommended to deploy modules as "unpacked" (exploded) WARs. This allows a faster startup and avoids unnecessarily increasing container size because Tomcat will not need to unpack WAR files.

The recommended way to add a module to an image is to:

- Copy the WAR to folder **/opt/ebx/webapps** . As stated previously, exploded format is recommended.
- Create an associated **Tomcat context XML file** named **module_war_name.xml** and copy it to folder **/opt/ebx/contexts**.
- Optionally, copy shared JARs to folder **/opt/ebx/lib** .

The **Tomcat context XML file** should have the following content:


```
<?xml version="1.0" encoding="UTF-8"?>

<Context docBase="${ebx.container.base}/webapps/module_war_name"/>
```

Using variable **`${ebx.container.base}`** is required for correct support of environment variable **`EBX_ROOT_PATH`**.

For more information on Tomcat contexts see https://tomcat.apache.org/tomcat-9.0-doc/config/context.html#Defining_a_context.

Here is a sample Docker file:

```
FROM ebx:6.1.0

COPY "./module_name.xml" "/opt/ebx/contexts"

COPY "./module_name" "/opt/ebx/webapps"
```

Adding a new locale

To add a new local you must have the jar file containing the language pack and add the locale to the ebx configuration.

Here is a sample Docker file

```
FROM ebx:6.1.0

COPY "<path_to_lib>" "/opt/ebx/lib"

RUN mv "/opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties" \

    "/opt/ebx/webapps/ebx/WEB-INF/ebx-default-original.properties"

RUN echo "ebx.file.previous=ebx-default-original.properties" >> \
```

```
"/opt/ebx/webapps/ebx/WEB-INF/ebx-default.properties"
```

```
RUN echo "ebx.locales.available=es" >> "/opt/ebx/webapps/ebx/WEB-INF/ebx-  
default.properties"
```

In this example, the lib is copied and the locale is set to "es" using the method described in the ["setting default configuration"](#) section

Adding a new JDBC driver

Adding a new JDBC driver is similar to adding a new library. You simply have to copy the jar file in the "/opt/ebx/lib" folder with the correct permission. Here is an example with the Oracle JDBC driver :

```
FROM ebx:6.1.0
```

```
ADD \
```

```
https://repo1.maven.org/maven2/com/oracle/database/jdbc/ojdbc11/21.8.0.0/  
ojdbc11-21.8.0.0.jar \
```

```
"/opt/ebx/lib/"
```

```
RUN chmod +r "/opt/ebx/lib/ojdbc11-21.8.0.0.jar"
```

See [Database drivers](#) for more information.

Technical administration

- [Repository administration](#)
- [UI administration](#)
- [UI – Workflow launcher](#)
- [Users and roles directory](#)
- [Data model administration](#)

- Database mapping administration
- Workflow management
- Task scheduler
- Audit trail
- Collaborative features (Team Up)
- Other

Repository administration

- Technical architecture
 - Overview
 - Rules for the database access and user privileges
 - Single JVM per repository
 - Failover with hot-standby
 - Repository status information and logs
- Auto-increments
- Repository management
 - Installation and upgrades
 - Inter-database migration
 - Repository backup
 - Archives directory
 - Repository attributes
- Monitoring management
 - Monitoring and cleanup of the relational database
 - Monitoring and clean up of the file system
- Dataspaces
 - Dataspaces/snapshots
 - Dataspace permissions
 - Repository history

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](#), accessed through a [configured data source](#).
- 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](#) and [Failover with hot-standby](#). Furthermore, to achieve horizontal scalability, an alternative is to deploy a [distributed data delivery \(D3\)](#) 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](#)
- [Supported databases](#)
- [Data source of the EBX® repository](#)

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.**

It is required for the database user specified by the [configured data source](#) to have the 'create/alter' privileges on tables, indexes and sequences. This allows for [automatic repository installation and upgrades](#).

See also:

- [SQL access to history](#)
- [Accessing a replica table using SQL](#)
- [Data source of the EBX® repository](#)

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. Moreover, is required for both servers to define the same value for `ebx.repository.directory`, and to share the directory defined by this value. (This is, in

particular, so that the Lucene indexes can be shared, i.e. not rebuilt on demand when the failover server starts.) 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](#) 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](#).

The format of the request URL must be:

```
http[s]://<host>[:<port>]/ebx?activationKeyFromStandbyMode={value}
```

- Using the Java API, call the method `RepositoryStatus.wakeFromStandby`.

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](#)' > 'Repository connection log'.

The status of the repository may be retrieved using the methods in the `RepositoryStatus` 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:

`state`

The state of the repository in terms of ownership registration.

- `D`: Java process is stopped.

	• O : Java process has exclusive ownership of the database. • S : Java process is started in failover standby mode, but is not yet allowed to interact with the repository. • N : Java process has tried to take ownership of the database but failed because another process is holding it.
heart_beat_count	The number of times that the repository has made contact since associating with the database.
info	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.

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:

Auto-increments	Lists all auto-increment fields in the repository.

Repository management

Installation and upgrades

Automatic installation and upgrades

By complying with the [Rules for the database access and user privileges](#), 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:

- The names of the database objects representing the mapped tables (history, replication) 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 `10000000.1234567890123456789` in Oracle will get rounded to `10000000.123456789012345679` 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](#)). This directory is automatically created during the first export from EBX®.

Attention:

As specified in the [security best practices](#), access to this directory must be carefully protected. Also, 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:

repositoryId	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> .
repository label	Provides a user-friendly label that identifies the purpose and context of the repository. For example: "Production environment".

store format	Identifies the underlying persistence system, including the current version of its structure.

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.

Database monitoring

The persistence data source of the repository must be monitored through RDBMS monitoring.

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).

History tables: impact on UI

For history tables, 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 dataspace and snapshots

In order to keep the cache and the repository to a reasonable size, it is recommended to close any dataspace and snapshots that are no longer required. This can be done in the following ways:

- Through the user interface, in the 'Dataspace' area.
- From the 'Dataspace / Snapshots' table under 'Dataspace' 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.closeHome`.
- Using the data service "close dataspace" and "close snapshot" operations. See [Closing a dataspace or snapshot](#) for more information.

Once the dataspace and snapshots have been closed, the data can be safely removed from the repository.

Note:

Closed dataspace and snapshots can be reopened in the 'Administration' area, under 'Dataspace'.

Deleting dataspace, snapshots, and history

Dataspace, 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 dataspace, snapshots, and history can be performed in a number of different ways:

- From the 'Dataspace/Snapshots' table under 'Dataspace' 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.deleteHome` and `RepositoryPurge.markHomeForHistoryPurge`.
- 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.purgeAll`.
- Using the task scheduler. See [Task scheduler](#) for more information.

The purge process is logged in the directory `${ebx.repository.directory}/db.purge/`.

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. This includes the dataspace and snapshots created for the persistent validation reports that have become obsolete. 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](#)

Monitoring and clean up of the 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, except for Lucene indexes:

- **Lucene indexes:** `${ebx.repository.directory}/indexes-(...)/`

Lucene indexes: Indexes can require a lot of disk space; they are critical to the correct functioning of EBX®. In nominal usage, they must not be deleted or modified directly. However, there are cases where deleting these indexes might be needed:

- If the repository is recreated from scratch, whereas the directory `${ebx.repository.directory}/` is preserved; to ensure consistency of data, it is then required to delete the root directory of the indexes.
- More generally, if the indexes have become inconsistent with the repository data (this could happen in rare cases of bugs).

After deletion, the content of the indexes will be lazily recomputed per table, derived from the content of the repository. The deletion must happen at the root folder of the indexes: if a single directory is deleted at a lower level, the global structure of the index will become inconsistent. This operation, however, has a cost, and should generally be avoided.

- **XML audit trail:** `${ebx.repository.directory}/History/`
- **Archives:** `${ebx.repository.directory}/archives/`
- **Logs:** [ebx.logs.directory](#)
- **Temporary directory:** [ebx.temp.directory](#)

Attention:

For **XML audit trail**, if large transactions are executed with full update details activated (contrary to the 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](#)
- [Tuning the EBX® repository](#)

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](#)

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 dataspace, snapshots, and history](#)

Dataspace permissions

This table lists all the existing permission rules defined on all the dataspace in the repository. You can view the permission rules and modify their information.

See also:

- [Dataspace permissions](#)

Repository history

The table 'Deleted dataspace/snapshots' lists all the dataspace that have already been purged from the repository.

From this section, it is also possible to delete the history of purged dataspace.

UI administration

TIBCO EBX® comes with a full user interface called [Advanced perspective](#) that includes all available features. The interface is fully [customizable](#) (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](#), giving the possibility to grant or restrict access to functional categories. Administrators can create simplified perspectives called [recommended perspectives](#) for end-users, containing only the features and menus they need for their daily tasks.

- [Global permissions](#)
- [Advanced perspective](#)
 - [Perspective creation](#)
 - [User interface](#)
 - [Child perspective menu](#)

- Recommended perspectives
 - Managing recommended perspectives
 - Resolution
- Custom views
 - Views
 - Views permissions
- User session management
- Accessibility mode
 - Supported screen reader

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.

Profile	Indicates on which profile the rule will be applied.
Restriction policy	Indicates if the permissions defined here restrict the ones defined for other profiles. See the Restriction policy concept for more information.
Dataspaces	Defines permissions for the Dataspaces area.
Data Models	Defines permissions for the Data Models area.
Workflow Models	Defines permissions for the Workflow Models area.
Data Workflows	Defines permissions for the Data Workflows area.

Data Services	Defines permissions for the Data Services area. Independently, it is also possible to: • Defines the access permissions for the REST built-in connector HTTP(S). This setting does not impact the REST Toolkit applications. • Defines the access permissions for the REST OpenAPI services. • Defines the access permissions for the SOAP connector HTTP(S) and JMS. • Defines the access permissions for the WSDL connector HTTP(S).
Administration	Defines permissions for the Administration area.

Note:

Permissions can be defined by administrators and by the dataspace or dataset owner.

Advanced perspective

The advanced perspective and its parameterization are unique. It is the parent perspective from which any [new perspective](#) 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](#)

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](#)
- [Favorite perspectives](#)

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](#)

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 \(deprecated\)](#). 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](#), 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:

User session default locale	Default session locale
Session time-out (in seconds)	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.

Login URL	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`,
- 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](#). After configuring the EBX® main configuration file, these configurations must be unset.

HTTP servlet policy	Header content of the servlet HTTP request: • if a field is not set, the default value in the environment configuration is used, • if a default value is not set, the value in the initial request is used.
HTTPS servlet policy	Header content of the servlet HTTPS request: • if a field is not set, the default value is chosen (in an environment configuration), • if a default value is not set, the value in the initial request is used.
HTTP external resources policy	Header content of the external resources URL in HTTP: • if a field is not set, the default value in the environment configuration is used, • if a default value is not set, the value in the initial request is used.
HTTPS external resources policy	Header content of the external resources URL in HTTPS: • if a field is not set, the default value in the environment configuration is used, • if a default value is not set, the value in the initial request is used.

Exit policy

Describes how the application is exited.

Normal redirection	Specifies the redirection URL used when exiting the session normally.
Error redirection	Specifies the redirection URL used when exiting the session due to an error. This feature is now deprecated and may be ignored by EBX®.
Redirection restrictions	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:

Activated	Indicates if the perspective is visible to authorized users.
Allowed profiles	The list of authorized user profiles for the perspective.
Allowed devices	The list of authorized devices for the perspective. If not specified, only "EBX® Web Application" can display this perspective.

Default selection	The menu item that is selected by default. This property is not available for the advanced perspective.

Application locking

EBX® availability status:

Availability status	This application can be closed to users during maintenance (but still remain open to administrators). Takes effect immediately.
Unavailability message	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.

IP access restriction	Restricts access to designated IP addresses (see IP pattern below).
IP restriction pattern	Regular expression representation of IP addresses authorized to access EBX®. For example, <code>((127\.\0\.\0\.\1) (192\.\168\.*\.*))</code> grants access to the local machine and the network IP range <code>192.168.*.*</code>.
Unique session control	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:

Max table columns to display	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.
Maximum auto-width for table columns	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.
Max expanded elements for a hierarchy	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.
Default table filter	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.
Display the message box automatically	Defines the message severity threshold for displaying the messages pop-up.
IE compatibility mode	Defines whether or not to compensate for Internet Explorer 8+ displaying EBX® in compatibility mode. 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'. See Specifying Document Compatibility Modes  for more information.
Forms: width of labels	The width of labels in forms.
Forms: width of inputs	The width of form input fields in forms.
Forms: height of text areas	The height of text entry fields in forms.
Forms: aggregated lists	The number of hidden candidate lines to be generated, available to create new instances in the list.
Forms: width of HTML editor	The width of HTML editors in forms.
Forms: height of HTML editor	The height of HTML editors in forms.
Searchable list selection page size	Maximum number of rows downloaded at each request of the searchable list selection (used for selecting foreign keys, enumerations, etc.).
Record form: rendering mode for nodes	Specifies how to display non-terminal nodes in record forms. If this property is modified, users are required to log out and log in for the new value to take effect.
Record form: display of selection and association nodes in creation mode	If enabled, the selection and association nodes will be displayed in record creation forms.

Display density	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.
Avatar displayed in the header	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.
Avatar displayed in the history	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.
Avatar displayed in the workflow	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

CSV file encoding	Specifies the default character encoding to use for CSV file import and export.
CSV : field separator	Specifies the default separator character to use for CSV file import and export.
CSV : list separator	Specifies the default list separator character to use for CSV file import and export.
Import mode	Specifies the default import mode.
Missing XML values as 'null'	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.

Web site icon URL (favicon)	Sets a custom favicon. The recommended format is ICO, which is compatible with Internet Explorer.
Logo URL (SVG)	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.

Logo URL (PNG)	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.
Logo URL (GIF/JPG)	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.
Main	Main user interface theme color, used for selections and highlights.
Header	Background color of the user interface header. By default, set to the same value as the Main color.
Workflow badge	Background and text/outline colors of new workflow task counters.
Primary buttons	Color of buttons selected by default. By default, set to the same value as the Main color.

Text of link style buttons	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.
Selected tab border	Border color of the selected tab. By default, set to the same value as the Main color.
Table history view: technical data	Background color of technical data cells in the table history view.
Table history view: creation	Background color of cells with the state 'creation' in the table history view.
Table history view: deletion	Background color of cells with the state 'deletion' in the table history view.
Table history view: update	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](#) 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](#)

Section Menu Item	This is a top level menu item. It contains other menu items.
Menu group	This is a container for other menu items.
Action Menu Item	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:

Type	The menu item type. See also: • Menu item types
Parent	The parent of the menu item. This property is not available for section menu items.
Label	The menu item label. 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.

Allowed devices	The list of authorized devices for this item. If not specified, all devices can display this menu item. Currently only two devices are supported: "EBX® Web Application" and "EBX® GO".
Icon	The icon for the menu item. Icon can be either "standard" (provided by EBX®) or an image, specified by a URL, that can be hosted on any web server. Icons size should be 16x16 pixels. This property is not available for section menu items.
Top separator	Indicates that the menu item section has a top separator. This property is only available for section menu items.
Action	The user service to execute when the user clicks on the menu item. See also: • User interface services 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. This property is only available for action menu items.

Selection on close	The menu item that will be selected when the service terminates. Built-in services use this property when the user clicks on the 'Close' button. This property is only available for action menu items.

Ergonomics and layout

See [Ergonomics and layout for the Advanced perspective](#) 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](#) 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.

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 [Perspective modes](#) 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](#) 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
```

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:

Documentation	Localized labels and descriptions.
Owner	Defines the user(s) owning and authoring this view definition.
View group	Indicates the menu group in which this view is displayed in the 'View' menu.

Share with	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
Recommend views	Allows the user to manage recommended views.	If the user is the dataset owner, the default value is 'Yes', otherwise it is 'No'.
Manage views	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'.
Share views	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'.
Publish views	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'.

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](#). 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.

Accessibility mode

An accessibility mode is available to help users with disabilities, when active it will tweak display to ensure that, wherever possible, the user interface is accessible.

The option can be found on the login screen, and also in the user pane once they are connected.

The selected value is persisted per user and shared among all its work sessions.

Allows to make all the controls of the form operable through a keyboard interface and thus conform to the WCAG Success Criterion 2.1.1.

Overview of changes applied when accessibility mode is on:

Tabular view	Add an "Access data" button to all row toolbars.
HTML editor	Make the toolbar being always expanded.
Forms	Expand/collapse buttons can now be focused.

Supported screen reader

The Jaws (version 2022) is the only screen reader supported by TIBCO EBX®. Any other technology may have some unexpected behavior.

UI – Workflow launcher

- [Introduction](#)
- [Workflow launcher in data section](#)
- [Creating and setting a launcher](#)
- [Activating workflow launcher](#)
- [Launching a workflow](#)
- [Adding a workflow launcher to the custom toolbar at table top](#)
 - [Adding a workflow launcher to a toolbar in other locations](#)

- [Access a launcher after workflow model modification](#)

Introduction

A **Workflow Launcher** is a **user service** for **launching workflows in TIBCO EBX®** directly from the data section without passing by the data workflow's inbox. This feature does not create workflow publications but launches existing ones. It offers several advantages, including the ability to launch workflow publication directly from the **data section** (table view, hierarchy view or record form view). In this way, the user experience is improved by avoiding to the user shifting his attention back and forth between the data section and the data workflow section. Hence, the user can launch a workflow while still focusing on his main task.

The second advantage is that it allows to launch the same workflow publication from any data selection. Thanks to the **dynamic mapping** of the [workflow data context](#) with the current data selection. The dynamic mapping offers the possibility to initialize the data context inputs at launch time. Hence, in order to launch the same workflow from n different data selections, it is not longer necessary to duplicate n times the same workflow model with different data selections or to provide an initial user service to configure the data to select.

The last solution is a programmatic solution which would solve the previously cited problems, however it is not the ideal solution because it does not fulfill the commitment zero line of code.

Understanding this feature requires familiarization with the following concepts:


CONCEPT	DESCRIPTION/LINK
Workflow model	See documentation
Data context	See documentation
Workflow publication	See documentation
Publication name	A unique identifier of a workflow publication
Data Workflow	See documentation
Data section	Is the data user interface which displays the datasets and tables in EBX. it is accessible from the main header.
Workflow launchers dataset	It contains two tables Launchers and Activations . With this dataset the user can configure the launchers of workflow

CONCEPT	DESCRIPTION/LINK
	publications and activate them for a particular user(s) and table(s). It is available in the Administration area in a section called workflow management .
Launcher	An entity which is used by the service workflow launcher to identify the workflow publication to launch and how to initialize its data context .
Launchers	A table in the workflow launchers dataset, each record is a Launcher type. In order to be able to launch a workflow publication from a data section, a launcher which points to a workflow publication should be added in this table.
Activation	An entity which hides or shows the launcher of a workflow publication for a given user profile(s) and for a particular table(s) from the data section.
Activations	A table in the workflow launchers dataset, each record is an Activation type. In order to make a workflow available on the toolbar of a particular table and for a particular user, a record with the corresponding launcher should be created in this table.
Dynamic mapping of data context	Is the process of initialization of data context input variables when launching a workflow publication. Before, the values of the data context were static defined at modeling phase or set dynamically with Java coding, but now it is possible to set dynamically the values of the data contexts, i.e. at launch time, with zero line of code.
Reserved variables for data context	A set of reserved variables. These variables define which data selection should be mapped with a data context variable. The possible values are: #{dataspace} , #{dataset} , #{xpath(table)} and #{xpath(record)} . For

CONCEPT	DESCRIPTION/LINK
	example, if the value <code>\${xpath(table)}</code> is used for a data context variable, this means that this variable will be mapped with the current adaptation table reference, at launch time.

Workflow launcher in data section

As previously stated, workflow launchers are now available directly on the toolbar of tables, table rows, records and hierarchies. There are a number of ways to display a workflow launcher on the toolbar of a table in data section. The display depends on the type of the toolbar (default toolbar or custom toolbar in the DMA) and whether or not [smart filtering](#) is applied. In the case of default toolbar, the action menu displays all the workflow launchers in a separate submenu called **Workflows** (see the screenshot below).



In the case of a custom toolbar, it is possible to define an action button (see the screenshot below) or action menu item for a particular workflow launcher. And finally, if the smart filtering policy is activated, then all the workflow launchers that are displayed using an action button or action menu item will not appear in the Action menu.

The screenshot shows the TIBCO EBX interface for 'Master Data: Students'. On the left is a sidebar with 'Notes', 'Actions', 'Courses', and 'Students' (selected). The main area displays a table of students. Above the table is an 'Actions' toolbar with a button 'Create a student' highlighted by a red box. The table has columns for 'Identifier' and 'First Name'.

	Identifier	First Name
☐	1	student1
☐	2	02
☐	3	03

In order to display a workflow launcher on a toolbar, first, a **launcher** must be **created** and **configured** (see [creating and setting workflow launcher](#) section), then, an **activation** should be **created** for this launcher (see [activating workflow launcher](#) section). Note that, if the current user has no [rights to launch a workflow publication](#) then the workflow launcher will not be available on the toolbar. Also, if the workflow publication is deleted or if there is any errors or warnings in the configuration or activation of the launcher, then the corresponding workflow launcher will not be available on the toolbar.

The title and tooltip of the button that will be displayed on the toolbar, are computed in the following order of priority: the custom documentation of the launcher activation is used, otherwise the documentation of the launcher is used. If description is left empty, then the following description is used "This user service will launch a data workflow.".

Note:

Particular attention should be paid for the workflow launchers which are **available on the record form and table rows**. Only the **launchers which requires record selection** are displayed on the toolbar of a record form and table rows. A workflow launcher which requires record selection is the one for which one of the data context variable is mapped with the reserved variable `${xpath(record)}` in the configuration of its launcher.

Creating and setting a launcher

A launcher is the entity which is used by the service Workflow Launcher to identify the workflow publication to launch and how to initialize its Data context. In order to create a launcher for a particular workflow publication, first navigate to the administration area, then the workflow management section, select the workflow launchers dataset, select the launchers table, and then add a new record. The second step, is to setup the following fields of the record:

FIELD NAME	DESCRIPTION
Launcher name	A unique identifier. It is used to select a launcher in the activation phase .
Workflow publication name	Defines the workflow publication that will be launched when the user service "workflow launcher" is executed. The workflow model should be first defined and published to be available in the list.
Data context variables	This field contains the list of input variables which were defined in the data context configuration phase of a workflow model. Each line is composed of a label, a value, and a toggle button to switch between default and overwritten value (see the screenshot below). The label is the name of the variable set by the user in the data context of the workflow model. By default, the value of each variable is the default one set in the data context. If the default value field is left empty in the data context, then the value of the variable is set to undefined. If the toggle button is set on overwritten value mode, then a wizard is made available on the right of the input field which allows to select a reserved variable.

FIELD NAME	DESCRIPTION
	When overwriting the value of a data context variable, two options are possibles: the override value may be a constant or a reserved variable. If a constant is used, then the value of the data context variable will not depend on the entity selection at launch time. However, if a reserved variable is used, such as `\${dataspace}`, the value of the corresponding variable will be mapped to the current entity selection, for example the data context variable which is assigned to `\${dataspace}` will be initialized with the current dataspace at launch time. Note: If there are no input variables in the data context of the workflow model, then this field should be hidden. If the workflow publication name field changes then this input is updated automatically and displays a new list of data context.
Documentation	A documentation is composed of a label and a description. The default value of the label is the publication name which is inherited from the field "workflow publication name" and the default value of the description is empty. If the "Publication name" field changes then the documentation will update automatically." The label and description can be overridden if needed thanks to the edit button on the right of the documentation widget." The value of the documentation will be inherited by the activation. Note that the documentation defined in the activation has

FIELD NAME	DESCRIPTION
	a higher priority over the one defined in the launcher.

At the bottom of a launcher record, a table of all activations of the current launcher is displayed (see the section Launcher activations in the screenshot below).

The screenshot displays the TIBCO EBX 'createRecord' form for a workflow launcher. The form includes the following elements:


- Launcher Name:** createRecord
- Workflow Publication Name:** testCP-13244
- Documentation:** English (United States) testCP-13244
- Data Context Variables:**
 - dataspace: \${dataspace}
 - dataset: \${dataset}
 - table: \${xpath(table)}
- Wizard and Toggle Button:** Indicated by red arrows pointing to icons on the right side of the Data Context Variables section.
- Launcher activations:** A section at the bottom showing a table of activations. The table has columns for Data Model, Documentation, Tables, and Profiles. The current activation is 'Publication: test', 'Create a record', 'table', and '[see details]'.
- Buttons:** Save, Save and close, Revert, and Close.

Activating workflow launcher

An activation is the entity which controls the availability of a workflow launcher on a table or record form. In order to display a workflow launcher in a toolbar of a particular table, first [create a launcher](#), then navigate to the administration area, select the workflow management section, select the workflow launchers dataset, select the activations table, and then add a new record. The second step, is to setup the following fields of the record:

FIELD NAME	DESCRIPTION
Launcher	A unique identifier of a launcher which is associated to a publication name of the target workflow publication . This field allows to select the launcher which will be displayed in the data section.

FIELD NAME	DESCRIPTION
Data model	The schema reference of a published Data Model. Only schemas that are published and used in Datasets are available.
Tables	The identifier of a table which will display the workflow launcher. It is possible to select one or many tables. This selector displays all tables which are contained in the field Data Model .
Profiles	List of profiles which are allowed to see and launch the workflow launcher in the previously selected tables.
Confirmation	This field allows displaying or not a dialog box to confirm launching of a data workflow (see screenshot about confirmation message . By default this feature is deactivated, in order to display the dialog box, the value "Enabled" should be checked (see screenshot below).
Documentation	A documentation is composed of a label and a description. The default value of the label and description are inherited from the documentation field of the launcher. This field is used to display the title of the button of the user service in the toolbar. The description is displayed when the user hovers over that button. The documentation of the activation has a higher priority over the one of the launcher. Note that, if the description of the documentation is left empty then the following one is displayed "This user service will launch a data workflow. In the case of the custom toolbar, the value of this field is also used for action button and action menu item if the documentation field of the action button or action menu item is left empty.



Launching a workflow

Once a launcher [created](#) and [activated](#) the corresponding workflow launcher becomes available on the toolbar's action menu. To launch the workflow click on the button of the workflow launchers(see [screenshot about default display of a workflow launcher](#)).

If the option [Confirmation](#) is enabled, then a dialog box displays to ask the user the confirmation before launching the selected data workflow (see [screenshot about data workflow launch confirmation](#)). Otherwise, the dialog box does not display and different outcomes are possible:

- If the option [automatically open the first step](#) is activated and the user has the [rights to execute the first work item](#) of the workflow, then the workflow launches and the first work item displays directly in the workspace of the Data section. While displaying the work item in the workspace, the origin data selection still displayed on the breadcrumb which allows to the user to maintain an overall contextual awareness (see [screenshot about displaying a work item on the data section](#)).
- If the option [automatically open the first step](#) is deactivated and the user has the [rights to launch the workflow](#), then the workflow is launched and added to the [data workflow inbox](#). An information message informs the user that the workflow has been launched successfully and that it is necessary to go to the [inbox of the data workflow section](#) to display and execute the first work item.
- If the workflow requires record selection (see [note in the section workflow launcher in data area](#)), the workflow is launched and displayed if and only if a record is selected otherwise an

error message is displayed informing the user that a record should be selected. A record can be selected manually from a table view or automatically when displaying a record form.

- If more than one record are selected, then a warning message notifies that only one record should be selected.

The top screenshot shows the TIBCO EBX 6.1.0 interface. The main window displays 'Master Data - Reference' with a dropdown menu set to 'table'. A modal dialog is open, asking 'Launch the Create a record ?' with 'OK' and 'Cancel' buttons. The background shows a table view with columns 'id', 'name', and 'firstName', and a message 'No records found.'

The bottom screenshot shows the same interface, but the 'table' dropdown is expanded, showing 'Step [0]' and 'New record'. The 'New record' form is displayed, with fields for 'id', 'name', and 'firstName'. The 'id' field is highlighted with a red box and labeled 'Source table'. The 'name' and 'firstName' fields are also highlighted with red boxes and labeled 'Work item'. A 'Save' button is visible at the bottom of the form.

Adding a workflow launcher to the custom toolbar at table top

The workflow launchers can be made available not only on default toolbar but also on [custom toolbar](#). A workflow launcher can be added on a custom toolbar as an [action button](#) or as an [action menu item](#) in a [custom menu group](#). Adding an action button or an action menu item to launch a workflow on a custom toolbar follows a similar procedure as for common user services: first, set the **target** field to **current context** otherwise the user service **workflow launcher** will not be available, then select the user service **workflow launcher**. When the service workflow launcher is selected, a new field **Launcher**, exclusive for workflow launcher user service, appears below the service input field (see [screenshot about adding a workflow launcher on custom toolbar](#)).


The field **Launcher** displays the list of launchers that have already been created in the table "Activation/workflow management/workflow launcher/Launchers" (see [section about creating and configuring launchers](#)) . Therefore, before adding a workflow launcher in a custom toolbar the launcher should be created and configured for this workflow launcher.

As for default toolbar, in order to show the workflow launcher on a custom toolbar its launcher should be activated in "Activation/workflow management/workflow launcher/activations". A quick link to access the activations table of the launchers is displayed on the right of the field **Launcher** (see [screenshot about adding a workflow launcher on custom toolbar](#))

The field Label and description displays the label and description inherited from the launcher (see [screenshot about adding a workflow launcher on custom toolbar](#)) , then if the launcher field changes, the label and description should update automatically. This field can be overridden to customize the label and description of the action button on the toolbar. The label and description that will be displayed on the toolbar, are computed in the following order of priority: the custom value of the field Label and description of the action button is used, if this field is left empty or contains the default label of the launcher, then the label and description of the launcher activation are displayed. If this last one is left empty, then the following description is used "This user service will launch a data workflow."

Note:

The label and description which are displayed on the toolbar in the data section and the label of the action button in the toolbar tree in the DMA could be different. This is the case when the field "Label and description" of the "action button" is left empty (neither default nor custom label is defined). The label of the "action button" in the toolbar tree in the DMA is inherited from the label of the launcher, however, the label and description on the toolbar in the data section are set to the label and description of the activation of the launcher.



Adding a workflow launcher to a toolbar in other locations


The workflow launcher service is only available on the toolbar that displays in the following locations: table view top, table rows, record form view top, or hierarchy view top. Until now, this service is not available on the toolbar that displays on top of the association table or in the association table rows.

Access a launcher after workflow model modification

After [creating a launcher](#) which points to a [workflow publication](#), the workflow model of this publication may be changed and republished. For instance, the data context of the workflow model may be changed: one or more variable data context can be added; one or more variable data context can be removed; the name of a variable can be changed. In this particular case, the launchers which points to this workflow publication should be reviewed and validated (saved) by the user because the data context used by these launchers is no more valid and it should be updated to match the one of the up-to-date workflow publication.

After republishing a workflow model, the user is notified if any of the workflow launcher points to the current workflow publication and if it should be reviewed. For that purpose, after publishing a workflow model, a preview button is displays allowing a quick access to the launchers in question . Note that if the user has no rights to access the workflow launchers list, then the names of those launchers are displayed .

When the user access to one of these launchers, via the preview button in the workflow model section or via the Administration section, the persisted data context is replaced with the up-to-date data context of the workflow publication, however the new data context is not yet persisted for the current launcher. In order to update the data context of the launcher, the user should first review the values and then save to accept the new changes (see [screenshot about accessing a launcher after modification of the data context of the corresponding workflow publication](#)).



Users and roles directory

- [Overview](#)
- [Concepts](#)
 - [Policy](#)
 - [Users](#)
 - [Roles](#)
- [Default directory](#)
 - [Directory content](#)
 - [Password recovery procedure](#)
 - [Administrator recovery procedure](#)
- [Custom directory](#)

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](#)
- [Custom directory](#)

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](#).

See also:

- [profile](#)
- [role](#)
- [user](#)
- [administrator](#)
- [user and roles directory](#)

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.

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](#).

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](#).

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](#):

```
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.

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®. The default login page can also be replaced by a specific company page.

See also:

- `DirectoryFactory`
- `Directory.getUserAuthenticationURI`

Data model administration

- [Administering publications and versions](#)
- [Migration of previous data models in the repository](#)
- [Schema evolutions](#)
 - [Impacts on data persistence](#)
 - [Impacts on side features](#)

Administering 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.

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.

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 link describes how the evolutions of data models are managed at the persistence level:

[Purging master tables in the database.](#)

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.

Database mapping administration

- [Overview](#)
- [Renaming columns in the database](#)
- [Purging columns in the database](#)
- [Renaming master tables in the database](#)
- [Renaming auxiliary tables in the database](#)
- [Purging master tables in the database](#)

Overview

Information and services relative to database mapping can be found in the *Administration* area.

See also:

- [Mapped modes](#)
- [DatabaseMapping](#)

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](#)

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
- Disabling replication on a specific field or group

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.

Renaming master tables in the database

This feature allows renaming master tables for history tables in the database. 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 for the repository prefix to be 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.

Renaming auxiliary tables in the database

This feature allows renaming history auxiliary tables in the database. This feature is not available for replicated tables since their names are specified in the data model.

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 for the new identifier to begin with a letter.

It is required for the new identifier to begin with the repository prefix.

It is also required for the repository prefix to be 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.

Purging master tables in the database

This feature allows purging history in the database if it is 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.

- Deactivate historization of the table in the data model, or
- Remove the table from the data model

Workflow management

- [Workflows](#)
 - [Execution of workflows](#)
 - [Workflow publications](#)
 - [Configuration](#)
- [Interactions](#)
- [Workflow history](#)
 - [Clean history](#)

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](#)

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](#)

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](#)

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](#)

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:

Cache expiry (seconds)	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.
User interface refresh periodicity (seconds)	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.

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.

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](#)

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](#).

Task scheduler

- [Overview](#)
- [Configuration from EBX®](#)
- [Cron expression](#)
 - [Format](#)
 - [Special characters](#)
 - [Examples](#)
- [Task definition](#)
- [Task configuration](#)

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.

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.

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 composed 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 composed 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.

EXPRESSION	MEANING
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).

Task definition

EBX® scheduler comes with some predefined tasks.

Custom scheduled tasks can be added by the means of [scheduler](#) Java API.

The declaration of schedules and tasks is done by selecting 'Task scheduler' in the 'Administration' area.

Task configuration

A user must be associated with a task definition; this user will be used to generate the [session](#) 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.util.Date`
- `java.net.URI`
- `java.net.URL`

Parameter values are set in XML format.

Audit trail

- [Overview](#)
- [Log content](#)
 - [Common header](#)
 - [Table view events](#)

- User service
- Legacy XML Audit trail (deprecated)

Overview

The audit trail feature allows to log all views access and user services launch in TIBCO EBX®.

The feature is disabled by default; to enable audit trail logging, configure the `audit` logging category with a level that enables `INFO` messages.

See also:

- [Activating the audit trail](#)

Log content

The information logged by audit trail comes in 2 parts: the common header and the event-specific content described in the following sections.

Common header

Here are the header attributes that are always logged:

user-login	Login of the user
user-ip	IP address of the user, when available
event	Type of event performed by the user, among the following value: • <code>TABULAR_ACCESS</code>: when a tabular view is accessed in the UI; • <code>HIERARCHY_EXPAND_NODE</code>: when a node is expanded in a hierarchical view, or when the view is displayed;

- **HIERARCHY_EXPAND_ALL**: when several nodes are expanded in a hierarchical view;
- **DATASPACE_PERMISSION_TABULAR_ACCESS**: when the permission table for a dataspace is accessed;
- **DATASET_PERMISSION_TABULAR_ACCESS**: when the permission table for a dataset is accessed;
- **USER_SERVICE**: when the export, compare, or validate built-in service, an addon or a custom user service is launched.

Table view events

This section describes the events specific to table views, which are:

- **TABULAR_ACCESS**
- **HIERARCHY_EXPAND_NODE**
- **HIERARCHY_EXPAND_ALL**
- **DATASPACE_PERMISSION_TABULAR_ACCESS**
- **DATASET_PERMISSION_TABULAR_ACCESS**

These events are triggered when the table view is opened or refreshed, when a search is performed and when the user navigates using pagination.

Here are the content attributes specific to these events:

dataspace	Key of the dataspace containing the table
dataset	Key of the dataset containing the table
table	Path of the table being accessed (does not apply to permission events)

view-id	Identifier of the custom view currently applied if any (does not apply to permission events)
view-name	Published name of the custom view currently applied if any (does not apply to permission events)
pagination-info	Information about the pagination of the tabular view or the expanded node
hierarchy-expanded-node	Identifier of the node being expanded (only applies to hierarchy events). When several levels are expanded (search or expand all), it is the node on the higher level. The identifier is the node label prefixed by ":". The root node is identified as ":". When a hierarchy is initially displayed, only the root node is expanded so the event is HIERARCHY_EXPAND_NODE.
filters	All filters applied to the view, aggregated in a JSON format. There are 3 kinds of filter: <code>quick-search</code>: user input in the quick search field, if any; <code>structured-search</code>: user advanced filters configured using the advanced search panel in a JSON format, if any; <code>view-filter</code>: filter defined by the view in a XPath format, if any. For instance: <pre>filters={"structured-search":["starts-with","/fi</pre>

User service

Here are the content attributes specific to this event:

service-key	Service key of the service being executed
service-mode	A service can be executed on any intermediate level of the selection entity. For example, if the current selection entity is a record, then the service may be executed on the following levels: dataspace, dataset, table or a record. This attribute allows to define on which level of the selection entity the service is executed.
selection	Entity selection on which the service is executed (branch or version, dataset, etc.) in a JSON format. When more than 100 records are selected, only the 100th first records are logged.
service-parameters	Additional input parameters specified for the service execution : • for built-in services, see the list of parameters; • for custom services, the parameters are specified in the declaration; • for addon services, see the addon documentation.

Legacy XML Audit trail (deprecated)

XML audit trail logs data updates to XML files. This legacy feature, now deprecated, will be removed in a future version. Although this type of audit trail should no longer be used, the [legacy documentation](#) is still available.

Collaborative features (Team Up)

- [Introduction](#)
- [Getting started with Team Up](#)
 - [Registering tables with Team Up](#)
 - [Managing user access and Team Up features](#)
- [Managing tags and tag catalogs](#)

Introduction

Enable Team Up to facilitate collaboration and engagement between business users. They can use Team Up to leave comments on records, rate records, and assign tags to records. Access to, and visibility of Team Up features is highly customizable. Global settings are available at the data model level, but can be overridden at the dataset level.

The following provides an overview of an administrator's role with Team Up features:

- **Comments:** Enable commenting on records. Assign moderators to tables that can manage comments from other user profiles.
- **Tags:** Enable tagging of records. Tags are classified in tag catalogs. Assign an administrator to manage tag catalogs.
- **Ratings:** Enable ratings for records.

Getting started with Team Up

If you have administrative access to an EBX® data model, register individual tables to enable Team Up. By default, when a table is registered, all Team Up features are set to active and visible; after the data model is published, users with Read-write or Read-only access can use Team Up. Use Team Up configuration settings to restrict user access and individual features. Additionally, you can configure visibility of each feature. For example, a feature can be active but hidden; users will not see them but previous comments, ratings, or tags are still associated with records.

Registering tables with Team Up

Use the **Team Up** data model extension to register each table:

1. Navigate to *Extensions > Team Up > Table configuration* and create a new record.
2. Use the **Table path** menu to select the table where you want to activate Team Up features.

3. Define user permissions, feature availability, moderators, and tag catalogs using the **Ratings**, **Comments**, and **Tags** groups.

See [Managing user access and Team Up features](#) for additional details.

4. After saving changes, publish the data model.

Managing user access and Team Up features

Manage user access to Team Up features at the data model and dataset levels. The table below shows the settings available in each location:

	In data models:	In datasets:
Activate features	Yes	No
Manage feature visibility	Yes	No
Restrict feature access permissions	Yes	Yes ***
Add moderators for comments	Yes	Yes
Assign tag catalog to table	Yes	No
Configuration location:		
	When viewing the data model navigate to: <i>Extensions > Team Up</i> and select the table you want to configure.	From the dataset's Actions menu select Team Up features configuration . Expand the features for each table you want to configure. *** Note: Permission restriction at the dataset level overrides data model permissions.

Managing tags and tag catalogs

Use tags and tag catalogs to help find and organize related data. Each table that has Team Up enabled is automatically associated with the default tag catalog. You can change the associated catalog. Only tags in the associated catalog are assignable to the table's records. The following table describes how to perform actions related to tags and tag catalogs:

Creating tag catalogs and administrators	• With administrative access, create tag catalogs in: <i>Administration > Metadata management > Team Up Administration > Tags > Tag catalogs</i>.

	Assign an administrator to a catalog: <i>Administration > Metadata management > Team Up Administration > Tags > Tags catalog administrators.</i>
Creating tags	- A new tag is automatically created when a business users assigns a tag to a record that does not exist in the tag catalog. - With administrative access, create tags in: <i>Administration > Metadata management > Team Up Administration > Tags > Tags.</i>
Assigning tags to catalogs	Change the catalog a tag belongs to: <i>Administration > Metadata management > Team Up Administration > Tags > Tags.</i> Note that this tag will no longer be available for records it was assigned to.

Other

- Lineage
- Event broker
 - Overview
 - Administration

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.

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

Event broker	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.
Topic	Corresponds to the EBX® event type that contains messages. The number of subscribers registered to a topic is unlimited.
Subscriber	Client implementation in the modules that receive the events related to the subscribed topic(s).

Topics

Dataspace and snapshot	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).
Repository	Corresponds to operations in the repository, such as: start-up and purge.
User session	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.

The event broker is based on a thread pool mechanism. The maximum number of threads can be defined in the properties file as follows:

```
# Defines the number of thread pool executors to  
  
# guarantee the publication of asynchronous events.  
  
# The default value is 2  
  
ebx.eventBroker.threadPool.size=2
```

Distributed Data Delivery (D3)

- [Introduction to D3](#)
- [D3 broadcasts and delivery dataspace](#)
- [D3 JMS Configuration](#)
- [D3 administration](#)

Introduction to D3

- [Overview](#)
 - [D3 architecture](#)
 - [Involving third-party systems](#)
 - [Protocols](#)
- [D3 terminology](#)
- [Known limitations](#)
 - [General limitations](#)

- Broadcast and delivery dataspace limitations
- Administration limitations


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 dataspaces 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 dataspaces.

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 \(D3\)](#)

D3 terminology

broadcast	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.
delivery dataspace	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 Data services .
delivery profile	A delivery profile is a logical name that groups one or more delivery dataspaces. Replica nodes subscribe to one or more delivery profiles.
cluster delivery mode	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.
federation delivery mode	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.
Primary node	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.
Replica node	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.
Hub node	An instance of EBX® acting as both a primary node and a replica node. Primary delivery dataspace and replica node delivery dataspace must be disjoint.

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.
- 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.

D3 broadcasts and delivery dataspaces

- Broadcast
 - Scope and contents of a broadcast
 - Performing a broadcast
 - Conditions
 - Persistence
 - Destination
- Replica node registration
 - Scope and contents
 - Performing an initialization
 - Conditions
- Accessing delivery dataspaces
 - Data services
 - Access restrictions
 - D3 broadcast Java API

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 'Primary node configuration'.

See also:

- Services on primary nodes

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 `D3NodeAsMaster`.

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 '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).
- 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

Destination

On 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 already exists on the first broadcast or full synchronization, it will be overridden.
- If an existing dataspace with the same identifier as the delivery one is detected outside of the `ebx-d3-reference`, an error will be raised.

See also:

- [Known limitations](#)

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 'Primary node configuration'. See [Supervision](#).

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](#).

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'](#)

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.getBroadcastVersion`.

D3 JMS Configuration

- [JMS for distributed data delivery \(D3\)](#)
 - [Common declarations on primary and replica nodes \(for shared queues\)](#)
 - [Optional declarations on primary nodes \(for replica-specific queues\)](#)
 - [Examples of JMS configuration](#)

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](#) 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](#) and [Distributed data delivery \(D3\)](#) 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
<code>jms/EBX_D3MasterQueue</code>	Weblogic: <code>EBX_D3MasterQueue</code> JBoss: <code>java:/jms/EBX_D3MasterQueue</code>	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 <code>JMSType</code> . Java type: <code>javax.jms.Queue</code>
<code>jms/EBX_D3ReplyQueue</code>	Weblogic: <code>EBX_D3ReplyQueue</code> JBoss: <code>java:/jms/EBX_D3ReplyQueue</code>	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 <code>JMSCorrelationID</code> . Java type: <code>javax.jms.Queue</code>

RESERVED RESOURCE NAME	DEFAULT JNDI NAME	DESCRIPTION
<code>jms/EBX_D3ArchiveQueue</code>	Weblogic: <code>EBX_D3ArchiveQueue</code> JBoss: <code>java:/jms/EBX_D3ArchiveQueue</code>	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 <code>JMSCorrelationID</code>. If the archive weight is higher than the threshold specified in the property <code>ebx.jms.d3.archiveMaxSizeInKB</code>, the archive will be divided into several sequences. Therefore, the consumption is filtered using the header fields <code>JMSXGroupID</code> and <code>JMSXGroupSeq</code> instead. Java type: <code>javax.jms.Queue</code>
<code>jms/EBX_D3CommunicationQueue</code>	WebLogic: <code>EBX_D3CommunicationQueue</code> JBoss: <code>java:/jms/EBX_D3CommunicationQueue</code>	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 <code>JMSType</code> which corresponds to the current repository ID. Java type: <code>javax.jms.Queue</code>

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
Primary-Replica nodes architecture	Between a primary node and two replica nodes with shared queues	Between a primary node and a replica node with replica-specific queues
Hub-Hub architecture	Between two hub nodes with shared queues	Between two hub nodes with replica-specific queues


Between a primary node and two replica nodes with shared queues




Between a primary node and a replica node with replica-specific queues



Between two hub nodes with shared queues



Between two hub nodes with replica-specific queues



D3 administration

- Quick start
 - Declare an existing dataspace on the primary node
 - Configure replica node for registration
- Configuring D3 nodes
 - Runtime configuration of primary and hub nodes through the user interface
 - Configuring primary, hub and replica nodes
 - Configuring the network protocol of a node

- Services on primary nodes
- Services on replica nodes
- Supervision
 - Primary node management console
 - Primary node supervision services
 - Replica node monitoring through the Java API
 - Primary node notification
 - Log supervision
 - Temporary files

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* and *replica* nodes included).

Update the `ebx.properties` *primary* 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 'Primary configuration' administration feature.
3. Add the dataspace to be broadcast to the 'Delivery dataspace' table, and declare the allowed profile.
4. Add the [delivery profile](#) to the 'Delivery profiles' table (it must correspond to a logical name) and declare the delivery mode. Possible values are: [cluster mode](#) or [federation mode](#).
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](#) 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](#) 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](#) 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 '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.

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 'Primary configuration' feature from the 'Administration' area, where you will find the following tables:

Delivery dataspace	Declarations of the dataspace that can be broadcast.
Delivery profiles	Profiles to which replica nodes can subscribe. The delivery mode must be defined for each delivery profile.

Delivery mapping	The association between delivery dataspaces 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](#)

Primary node

In order to act as a *primary* node, an instance of EBX® must declare the following property 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
```

See also:

- [primary node](#)

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 property 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

#####
```

```
# 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](#)

Replica node

In order to act as a *replica* node, an instance of EBX® must declare the following property 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](#)

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](#)

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 \(D3\)](#)

Services on primary nodes

Services to manage a primary node are available in the 'Administration' area of the replica node under 'Primary node configuration' and also in the 'Delivery dataspaces' and 'Registered replica nodes' tables. The services are:

Relaunch replays	Immediately relaunch all replays for waiting federation deliveries.
Delete replica node delivery dataspace	Delete the delivery dataspace on chosen replica nodes and/or unregister it from the configuration of the D3 primary node. To access the service, select a delivery dataspace from the 'Delivery dataspaces' table on the primary node, then launch the wizard.

Fully resynchronize	Broadcast the full content of the last broadcast snapshot to the registered replica nodes.
Subscribe a replica node	Subscribe a set of selected replica nodes.
Deactivate replica nodes	Remove the selected replica nodes from the broadcast scope and switch their states to 'Unavailable'. Note: The "in progress" broadcast contexts are rolled back.
Unregister replica nodes	Disconnects the selected replica nodes from the primary node. Note: The "in progress" broadcast contexts are rolled back.

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 *Configuration of replica* node to manage its subscription to the primary node and perform other actions:

Register replica node	Re-subscribes the replica node to the primary node if it has been unregistered.
Unregister replica node	Disconnects the replica node from the primary node.

	Note: The "in progress" broadcast contexts are rolled back.
Close and delete snapshots	Clean up a replica node delivery dataspace. 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. Note: The last broadcast snapshot is automatically excluded from the selection.

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 'Primary node configuration'. They are as follows:

Registered replica nodes	Replica nodes registered with the primary node. From this table, several services are available on each record.
Broadcast history	History of broadcast operations that have taken place.
Replica node registration log	History of initialization operations that have taken place.
Detailed history	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 ' Primary node configuration'. The services are as follows:

Check replica node information	Lists the replica nodes and related information, such as the replica node's state, associated delivery profiles, and delivered snapshots.
Clear history content	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 ' Replica node configuration'.

See also:

- `NodeMonitoring`

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 '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](#)

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
```

Staging

- [Overview](#)
- [Glossary](#)
 - [Staging server](#)
 - [Staging element](#)
 - [Staging domain](#)
 - [Staging archive](#)
 - [Staging operations](#)
- [Performing staging](#)
 - [Specifying import behavior](#)
- [Staging domains and best practices](#)
- [Staging permissions and administrative delegation](#)
- [Error handling](#)
 - [Error handling during import operation](#)
 - [Error handling during export operation](#)
- [Supported types of staging elements](#)
 - [Data tables](#)
 - [Profiles](#)

- Data models
- Datasets
- REST API
- Temporary folders handling

Overview

The staging feature allows you to move TIBCO EBX® repository elements from one environment to another. For example, it can stage configuration settings from a development environment to a testing environment. Many elements are supported as part of the staging feature. However, its main purpose is to handle configurable elements contained in the repository and *not* data that are updated by end users. See [Supported types of staging elements](#) for a list of all supported types of staging elements.

By default, the staging features are active. Use the `ebx.staging.activated` parameter contained in the `ebx.properties` file to activate or to deactivate staging in different environments.

Note:

Administrative access, or delegated staging access is required to use the staging feature. See [Staging permissions and administrative delegation](#) for more information.

Glossary

Staging server

An environment composed of deployed software (including EBX®) that is dedicated to development, testing, or production. Staging is used to build, test, and review an updated, or new solution before it is moved (staged) from one server to another.

Staging element

A configuration element in an EBX® repository that is stageable. Different types of staging elements include: data model publications, workflow model publications, dataspace, etc.

Staging domain

A set of staging elements that can be transferred between EBX® servers.

Staging archive

File used to stage elements of a domain from a server to another.

Staging operations

The actions of importing and exporting a staging archive to and from a server.

Performing staging

The process of staging begins with the creation of domains to organize and stage elements. Once a domain is ready, it is exported as an archive file. Finally, the archive is imported into the target environment.

This operation can be performed in the EBX® graphical user interface, or by using the [REST APIs](#). See [Staging domains and best practices](#) for staging domain best practices.

Attention:

Only product versions with the same minor number are compatible for staging.

The following steps show how to use the user interface to perform a staging operation:

1. Create a domain:

See also:

- [Staging domains and best practices](#)

- a. In the source repository, navigate to *Administration > Repository management* and select **Staging**.
- b. Select **Domains** and click **Add a domain**.
- c. Provide a name and save, but do not close.

- d. Define the elements to stage by selecting the desired elements and if required, the specific content in the category. The list of selections is automatically saved on changes.

Note:

Any elements that are in a state of error are hidden and cannot be selected.

Note:

Additionally, each node in the repository is limited to displaying 1000 elements. It is possible to change this default value using the `ebx.staging.maxElementPerList` property contained in the `ebx.properties` file.

See also:

- [Activating and configuring staging](#)

2. Export and download the staging archive by selecting **Download** at the top of the screen or at each line of the domains table.
3. Import the staging archive:
 - a. In the target repository, navigate to *Administration > Repository management* and select **Staging**.
 - b. Click **Import a staging archive** at the top of the screen.
 - c. Browse to select the staging archive.
 - d. Choose one of the import options. See [Specifying import behavior](#) for more information.
 - e. Click **Import archive**.

See [Error handling](#) for more information.

Specifying import behavior

The following image and table highlight available options when importing a staging archive:

TIBCO® EBX®

aa

Administration

Staging ▾

Actions ▾

Domains

▼ Reports

Import operations

Export operations

▸ Configuration

Import a staging archive

Upload and import a staging archive into the current repository.

Choose file

No file chosen

☐ **Force archive import.**

Imports the archive, even if it is identical to the most recently imported archive.

☐ **Force import of components.**

Import each component even if the identical component exists in the repository.

☐ **Save the staging domain definition.**

The staging domain definition will be saved in the staging management. (Should be used only when the target repository is for development).

☐ **Dry run (testing).**

Run a staging import for testing purposes without modifying the repository.

Import archive

Force archive import

If the same staging archive was already successfully imported, the staging import is ignored by default. Use this option to force import execution of the archive. Use this mode if an initial import did not go as expected, then retry after making changes in the source repository.

Force import of components

If the same staging component was already successfully imported, it is ignored by default. Use this option to force import execution of the component. Use this mode after modifying an archive and force replay the entire persistence process before going into production.

Note:

601

	If the same archive has already been successfully imported, the option <code>Force archive import</code> must also be checked to force a component import. Otherwise, the entire import process will be skipped.
Save the staging domain definition	By default, staging domain records are not created in the Domains table when executing an import operation. Use this option to import manually the domain definition.
Dry run (testing)	This option is used for testing purposes to run an import without modifying the repository. Use this option to ensure the existence of declared dependencies.

Staging domains and best practices

Staging domains are an effective means to organize the repository configuration settings. Use a domain to break large repository configurations into smaller parts. This allows each part to have different life-cycles. For example, HR and Procurement applications require updates at different times. In this case, create a separate domain for each application. Then define and stage these domains in multiple environments (dev, test, prod, etc.). This method also avoids having to test the entire configuration repository each time.

Domains require a unique name that conforms to dot notation. Supported characters include alpha and numeric from ASCII, underscores, dashes, and dots (the name cannot begin with an underscore). This behavior can be used like a namespace to create a hierarchy. For example:

```
com.acme.hr , com.acme.procurement , com.acme.procurement.catalog ,
com.acme.procurement.catalog.2022 .
```

Some best practices include:

- Begin to clearly define domains as soon as possible in the staging life-cycle.
- Use a separate domain for permissions because development environment permissions are usually different from production permissions.

- Staging should only handle configuration data in the repository. The types "Data tables" and "Filtered data" are supported, but it should apply to data that are not updated by users (regions, countries, etc.) as the staging data will update them and might override changes by users.
- Dataspaces that include long-duration transactional data hold operational data, and not configuration data. They can be numerous and clutter up view of staging components. Prevent them from displaying using the settings available in: *Administration > Repository management > Staging > Configuration > Hidden dataspace*s.
- Create a backup of the database before importing one or more staging archives.

Staging permissions and administrative delegation

By default, administrators in EBX® have permission to access staging and can grant access rights to non-administrative users. Access rights can specify availability of individual actions, or all actions. However, since staging has the ability to create and update configuration elements, it is important to use caution when granting access to staging.

Grant a user staging administration permissions using administrative delegation. This process is completed using User permissions apply when selecting the content of a staging domain; they cannot select elements to stage that they don't have access to.

You can configure administrative delegation in the **Administration** panel under *User interface > Global permissions*.

See also:

- [Configuring the user and roles directory](#)

The following table highlights the permission differences between an administrator and a user with staging administration delegation:

Administrator	• Can export a staging archive. • Can import a staging archive.
A user with delegated staging administration	• Can export a staging archive.

- **Cannot** import a staging archive.

Note:

This type of user has access to, and can export domains created by other users or administrators even if the user does not have access to the staging elements. Choose this option when you want to ensure the consistency of the domain definition. Only users with write permissions can create domains.

Error handling

Error handling during import operation

An import operation may return errors and warnings. The import can be successful, partially successful, on failure, have elements that were ignored or totally ignored. The import operation does not stop on first error, but tries all components. When an error occurs, the component will not import correctly. It might require a change in either the source or target repository to complete successfully. The error message should provide enough information to help discern the issue.

Note:

Importing is not fully transactional and errors might lead to an unstable repository.

A dependency check is done, and if a required element is missing, an error is returned. This can happen when importing a domain that contains a dataset and the dataspace associated with the dataset is not in the domain, or in the destination repository.

When a warning occurs, the component has probably been imported correctly but might work in a degraded way, or is not accessible. During the dependency check, if a referenced optional element is missing, a warning is returned. This is the case of the owner profile of a dataspace.

Note:

Staging tries to avoid writing as much as possible by default. This behavior can be overridden to force create or update staging elements. See [Specifying import behavior](#) for more information.

Import report

An import report is generated on each import operation performed through the UI or REST. It includes the import status of the domain and its components. The report also contains messages with severity levels (`info` , `warning` , `error`) that can help troubleshoot issues.

Import domain status values:

- **Success:** All components are imported successfully. The import might ignore some components if they are already present and as expected in the destination repository.
- **Partially failed:** Some component imports raised errors. Other components were imported successfully. The [Status of each component](#) import should be checked.
- **Failed:** An exception occurred while reading the archive, or the component imports are all in error. Unlike the "Partially failed" status, no changes are performed to the repository.
- **Ignored:** The archive is ignored as it was already imported or all the component imports were ignored. No changes are performed to the repository.
- **Waiting:** Used in asynchronous mode. The import operation is queued and waiting to be processed.
- **Started:** The import process is started and not finished.

Note:

On system start, all imports that were started or waiting will pass automatically to failed because the job was interrupted.

Import component status values:

- **Success:** The staging element was successfully created or updated.
- **Failed:** An error occurred while importing the staging element.
- **Ignored:** The import is ignored since the element is as expected in the repository.
- **Started:** The import is started and not finished.

Error handling during export operation

Unlike import operations, the export operation only succeeds on full archive export; it fails if any error occurs. There is no partial states for staging exports.

If components in the domain are missing, an error is returned and the archive is not generated.

Export report

An export report is a report generated on each export operation performed through the UI or REST. It includes the export status of the archive.

Export domain status values:

- **Exported:** The archive was successfully exported.
- **Requested:** Used in asynchronous mode. The export is requested and currently processing.
- **Aborted:** The archive export was aborted due to export errors or a processing interruption.
- **Expired:** Using the REST service's asynchronous mode, an archive was generated but expired. It is no longer available.

Supported types of staging elements

The following types of repository elements are supported for the staging feature:

- Dataspace
 - Dataspace permissions
 - Dataset (see [Datasets](#) for more information)
 - Dataset permissions
 - Data table
 - Filtered data
- Data model (see [Data models](#) for more information)
 - Data model permissions
- Data model publication (see [Data models](#) for more information)
- Workflow publication
- Perspective

- Perspective permissions
- Profile (Roles) (see [Profiles](#) for more information)
- View publication
 - View publication permissions
- Administration
 - Recommended perspective
 - Recommended perspective permissions
 - Historization profile
 - History and logs permissions
 - Task scheduler permissions
 - Workflow execution permissions
 - Workflow history permissions
 - Workflow interaction permissions
 - Workflow launcher permissions
 - Workflow technical configuration
 - Dataspace administration permissions
 - Data modeling permissions
 - Auto increment permissions

Data tables

Data tables and filtered data are imported with a default configuration:

- The default [import mode](#) is update or insert
- The [blocking constraints](#) are disabled
- The [by delta](#) mode is enabled
- The [table triggers](#) are activated on import

Profiles

Only roles can be staged. Profiles are ignored when importing a staging archive, except for roles and if the directory is the default directory. When exporting a staging domain, a staging element

linked to one or more profiles (user or role) will be staged with the link(s) whenever the users or roles are in the directory or not.

Data models

Data model publications and data models that are located in modules are not available for staging. However, a data model can be staged even if its corresponding publication is located in a module (the publication will not be staged).

Datasets

It's not possible to stage child datasets. They don't display when adding staging element to a domain.

REST API

See [Staging operations](#) for more information.

Temporary folders handling

Temporary folders are created to support staging operations and asynchronous process. Each folder created expires after a configured delay. The scheduled task *DeleteExpiredStagingTemporaryFolders* ensures periodic deletion of expired folders.

Note:

The dedicated size on disk of temporary folders created for staging can be configured through the property `ebx.staging.maxTemporaryFolderSizeInBytes`.

Security Guide

- [Security Best Practices](#)

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. While these practices are commonly regarded as beneficial, they may not be relevant to your particular infrastructure and security policy.

- [Encryption algorithms](#)
- [HTTPS](#)
- [Installation](#)
- [Web Server](#)
- [Application Server](#)
- [Java](#)
- [Database](#)
- [Archive directory](#)
- [User directory and Administration rights](#)
- [Permissions](#)

Encryption algorithms

The Web Server or Application Server may specify encryption algorithms when setting HTTPS parameters. Some recommendations on these algorithms are provided in section [HTTPS](#). Password and fields having `osd:password` as a type store a hash of their value, using the SHA_512 algorithm. This is, notably, the case for the password of users of the default directory.

HTTPS

It is recommended to use HTTPS for communication with clients (GUI and REST or SOAP). 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 is 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, being too restrictive on allowed cyphers may prevent some clients from connecting, as they may not be able to negotiate a 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

Installation

Deployed components, such as 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.

Web Server

If you have to expose web applications on the internet, it is a good practice to protect them with a Web Server in a [demilitarized zone](#), while EBX® and the database server can be in a production zone. Consider the following practices for your configuration.

The secure cipher suite and protocols should be set according to the above section [HTTPS](#).

Do not use the default configuration, and remove any banner that might 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
```

Use the Web Server to set restrictions with HTTP security headers. Note that headers related to the origin 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, configure EBX® to not set any HTTP security headers. To do so, set the property `ebx.security.headers.activated` to `false` or unset.

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
```

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"
```

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](#).

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.

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.

EBX® allows a very high level of customization through custom code. All integrated Java modules are considered by EBX® as trusted. Hence, all development on top of EBX® should be reviewed and validated. As an example, developers should not generate HTML from values coming from the database without proper escaping. For more details on this, see the [Cross Site Scripting prevention on the OWASP site](#). Here is a proper escaping example: the name of a store is encoded before being displayed in an HTML form. The `StringEscapeUtils` class included in Apache Commons Lang is used for string encoding.

```
public class StoreMainPane implements UIFormPane

{

    public static final String STORE_NAME_STYLE = "font-weight: bold;
padding-top:20px; padding-bottom:20px";

    @Override

    public void writePane(final UIFormPaneWriter writer, final
UIFormContext context)

    {

        String storeName = (String)
context.getValueContext().getValue(Paths._Store._Name);

        writer.add("<div ").addSafeAttribute("style",
STORE_NAME_STYLE).add(">");
```

```
        writer.addSafeInnerHTML("Data stored for " + storeName);

        writer.add("</div>");

        // ...

    }

}
```

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.

Archive directory

On the server, the [archive directory](#) must be properly secured and/or encrypted. Indeed, any archive exported from the EBX® instance will be created there, and these archives are neither encrypted nor protected by password. As a consequence, any user with an access to these files will be able to see the content regardless of any permission defined in EBX®.

User directory and Administration rights

For production and test platforms, EBX® must be integrated with a [custom directory](#) 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.

Permissions

Special care is required when defining permissions in EBX®. Persons in charge of this are expected to be aware of the content of the [permission documentation](#), and especially the information provided in the [Important considerations about permissions](#) section.

Developer Guide

- Introduction
 - Packaging TIBCO EBX® modules
 - Mapping to Java
 - Tools for Java developers
 - Terminology changes
- Data model
 - Introduction
 - Data types
 - Tables and relationships
 - Constraints
 - Triggers and functions
 - Labels and messages
 - Additional properties
 - Data services
 - Toolbars
 - Custom forms
- Workflow model
- User interface
 - Interface customization
 - User services
 - Overview
 - Quick start
 - Implementing a user service
 - Declaring a user service
 - Development recommendations

- SOAP data services
 - Introduction
 - WSDL generation
 - SOAP operations
- REST data services
 - Introduction
 - Built-in RESTful services
 - JSON Formats
 - Introduction
 - Extended
 - Compact
 - Common
 - Others
- SQL in EBX®
 - Introduction
 - Comparison operators
 - Arithmetic operators and functions
 - Logical operators
 - String operators and functions
 - Date and time functions
 - EBX® SQL functions
- REST Toolkit
- EBX® Scripting
 - Record permission
 - Function field
 - Function field API
 - Unit summary
 - Unit default

- [Unit core.complex](#)
- [Unit core.date](#)
- [Unit core.datetime](#)
- [Unit core.list](#)
- [Unit core.locale](#)
- [Unit core.log](#)
- [Unit core.math](#)
- [Unit core.resource](#)
- [Unit core.string](#)
- [Unit core.time](#)
- [Unit core.uri](#)
- [Java APIopen JavaAPI in popup](#)

Introduction

- [Packaging TIBCO EBX® modules](#)
- [Mapping to Java](#)
- [Tools for Java developers](#)
- [Terminology changes](#)

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.

- [Module structure](#)
- [Module declaration](#)
- [Module registration](#)
 - [Deployment descriptor example](#)
 - [Registration example](#)

- Packaged resources
 - Directory structure
 - Example

Module structure

An EBX® module contains the following files:

<code>/WEB-INF/ebx/module.xml</code>	This mandatory document defines the main properties and services of the module. See Module declaration .
<code>/WEB-INF/web.xml</code>	This is the standard Java EE deployment descriptor. It can perform the registration of the EBX® module when the application server is launched. See Module registration .
<code>/META-INF/MANIFEST.MF</code>	Optional. If present, EBX® reports the 'Implementation-Title' and 'Implementation-Version' values to <i>Administration > Technical configuration > Modules and data models</i> .
<code>/www/</code>	This optional directory contains all packaged resources, which are accessible via public URL. See Packaged resources .

Required files for Oracle WebLogic server:

<code>/WEB-INF/weblogic.xml</code>	WebLogic deployment descriptor file which activates the <code>prefer-web-inf-classes</code> policy, such as the following: <pre><?xml version="1.0" encoding="UTF-8"?> <weblogic-web-app xmlns="http://xmlns.oracle.com <container-descriptor> <prefer-web-inf-classes>true</prefer-web-inf</pre>

```
</container-descriptor>
```

```
</weblogic-web-app>
```

See [weblogic.xml Deployment Descriptor Elements](#) for more information.

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
<code>name</code>	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.
<code>publicPath</code>	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	No.

ELEMENT	DESCRIPTION	REQUIRED
	of the module when computing absolute URLs. If this field is not defined, the public path is the module's <code>name</code> , defined above.	
<code>services</code>	Declares user services using the legacy API. See Declaration and configuration of legacy user services . From the version 5.8.0, it is strongly advised to use the new user services .	No.
<code>beans</code>	Declares reusable Java bean components. See the workflow package .	No.
<code>ajaxComponents</code>	Declares Ajax components. See Declaring an Ajax component in a module in the Java API.	No.

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 `ModuleRegistrationListener`.

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`: [JavaTM Servlet 3.0 Specification](#) for more information.

or:

- contain a Servlet extending the class `ModuleRegistrationServlet` ;
- 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 `ModuleRegistrationServlet` 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](#).
- 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](#).

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://java.sun.com/xml/ns/javaee"

    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

    xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
```

```

                                https://www.oracle.com/webfolder/technetwork/
jsc/xml/ns/javaee/web-app_3_0.xsd"

    version="3.0">

    <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 RuntimeException

    {

        // 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

            ...

        }

    }

```

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`
- [Type](#) `osd:resource`

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](#).

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

Mapping to Java

- [How to access data from Java?](#)
 - [Read access](#)
 - [Write access](#)
 - [Modification of mutable objects](#)
- [Transactions and concurrency](#)
 - [Concurrency](#)
 - [Query snapshot isolation](#)
 - [Adaptation objects](#)
- [Mapping of data types](#)
 - [Simple data types](#)
 - [Complex data types](#)
- [Java bindings](#)
 - [Benefits](#)
 - [XML declaration](#)

How to access data from Java?

Read access

Data can be read from various generic Java classes, mainly `Adaptation` and `ValueContext`. The getter methods for these classes return objects that are typed according to the mapping rules described in the section [Mapping of data types](#).

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 `ValueContextForUpdate`, or
- During the user input validation, by calling the method `setNewValue` of the class `ValueContextForInputValidation`.

Modification of mutable objects

According to the mapping that is described in the [Mapping of data types](#) 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.

Transactions and concurrency

Concurrency

At the dataspace level	In a single dataspace, the system supports running only one single read-write <code>Procedure</code> and multiple concurrent <code>ReadOnlyProcedure</code> s. Concurrent accesses outside any <code>Procedure</code> are also supported.
At the repository level	At the repository level, concurrency is limited for only some specific operations. For example (non-exhaustive list): • A data model publication excludes many operations. • A dataspace merge excludes write operations on the two dataspace involved in the merge.

Query snapshot isolation

The following table defines the properties related to query isolation. Note that a *query result* is represented in the Java API by either a `QueryResult` or a `RequestResult`:

Queries outside of a <code>Procedure</code>	Data is frozen at the time of fetching the query result. More precisely, a query result accesses only committed data as of the last committed transaction at the time of fetching this result. The content of this result never changes afterwards. A query outside of a <code>Procedure</code> can be considered as a self-containing <code>ReadOnlyProcedure</code>.
Queries inside of a <code>Procedure</code> , in the same dataspace as the <code>Procedure</code> underlying dataspace	The query result reflects the last committed state before the <code>Procedure</code> starts and the changes that occurred in the <code>Procedure</code> previously to the query result fetch. The content of this result never changes afterwards, whatever happens in the <code>Procedure</code>.
Queries inside of a <code>Procedure</code> , in another dataspace	The consistency is guaranteed at the repository level, so the query result reflects the last committed state before the <code>Procedure</code> starts. The content of this result never changes after the query is fetched, whatever happens in the whole repository.

Adaptation objects

In Java, a persistent dataset or a persistent record are both represented by an instance of the `Adaptation` class.

The following table defines the properties related to `Adaptation` objects.

Immutability	An <code>Adaptation</code> object instance is immutable.

	Therefore, the client code should not "hold" an <code>Adaptation</code> object for a long time (in particular beyond a transaction boundaries). However, it is possible to invoke the method <code>Adaptation.getUpToDateInstance</code>.
Fetch	If an <code>Adaptation</code> is fetched from a <code>QueryResult</code>, then the snapshot isolation rules described in the previous section apply. Otherwise, if an <code>Adaptation</code> is fetched from a running <code>Procedure</code>, it reflects the last committed state before the Procedure starts. Otherwise, outside of a <code>QueryResult</code> or a running <code>Procedure</code>, the <code>Adaptation</code> reflects the state of the record on its fetch-time. See also: • <code>AdaptationHome.findAdaptationOrNull</code> • <code>AdaptationTable.lookupAdaptationByPrimaryKey</code>

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](#).

See also:

- `SchemaNode.createNewOccurrence`

Multiple cardinality on a simple element

If the attribute `maxOccurs` is greater than 1, then 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 must access these values, use a custom JavaBean. To do so, use the `osd:class` declaration described in the next section.

You can transparently instantiate, read and modify the mapped Java object, with or without the attribute `osd:class`, by invoking the methods `SchemaNode.createNewOccurrence`, `SchemaNode.executeRead` and `SchemaNode.executeWrite`.

Mapping of complex types to custom JavaBeans

You can 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`.

```

<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](#), or
- An instance of `AdaptationTable`, if the property `osd:table` is specified.

Java bindings

Java bindings support 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](#).

Benefits

Ensuring the link between XML Schema structure and Java code provides a number of benefits:

- **Development assistance:** Auto-completion when you type an access path to parameters, if it is 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, you can easily verify where a parameter is used.

Consequently, it is strongly encouraged that you 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/">
        <javaPathConstants typeName="com.creditonline.RulesPaths">
          <nodes root="/rules" prefix="" />
        </javaPathConstants>
```



```

        <javaPathConstants
typeName="com.creditonline.StylesheetConstants">

            <nodes root="/stylesheet" prefix="" />

        </javaPathConstants>

    </ebxnd: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`.

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.

- [Activating the development tools](#)
- [Data model refresh tool](#)
- [Generating Java bindings](#)
- [Path to a node](#)
- [Web component link generator](#)

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` using the property `backend.mode=development`.

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.

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](#)

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.

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.

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](#)

- [Terminology changes in version 5.9](#)
- [Terminology changes in version 5.0](#)

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

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
Dataset	Adaptation instance
Child dataset	Child adaptation instance
Data model	Data model

NEW TERM	TERM PRIOR TO VERSION 5.0.0
Dataspace	Branch
Snapshot	Version
Dataspace or snapshot	Home
Data Workflow	Workflow instance
Workflow model	Workflow definition
Workflow publication	Workflow
Field	Attribute
Inherited field	Inherited attribute
Record	Record/occurrence
Validation rule	Constraint
Simple/advanced control	Simple/advanced constraint

Data model

- [Introduction](#)
- [Data types](#)
- [Tables and relationships](#)
- [Constraints](#)
- [Triggers and functions](#)
- [Labels and messages](#)
- [Additional properties](#)
- [Data services](#)
- [Toolbars](#)
- [Custom forms](#)

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](#), such as integer, boolean, decimal, date, time;
- The ability to define additional [simple types](#) and [complex types](#);
- The ability to define simple lists of items, called [aggregated lists](#);
- [Validation constraints](#) (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](#), for example: locale, resource, html;
- Definition of [tables](#) and [foreign key constraints](#);
- Mapping data in EBX® to Java beans;
- [Advanced validation constraints](#) (extended facets), such as dynamic enumerations;
- Extensive [presentation information](#), 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.

- [Editing the data model](#)
- [References](#)
- [Relationship between datasets and data models](#)
- [Pre-requisite for XML Schemas](#)
- [Conventions](#)
- [Schemas with reserved names](#)

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](#).

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.

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](#)

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](#)

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>

```

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

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.

Data types

This chapter details the data types supported by TIBCO EBX®.

See also:

- [Tables and relationships](#)
- [XML Schema built-in simple types](#)
- [XML Schema named simple types](#)
- [XML Schema complex types](#)
- [Extended simple types defined by EBX®](#)
- [Complex types defined by EBX®](#)
 - [Restrictions on User Defined Attributes and Catalogs](#)
- [Aggregated lists](#)
- [Including external data models](#)

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
<code>xs:string</code>	<code>java.lang.String</code>	The default strategies defined for string fields are detailed in the section Default strategy for string fields .
<code>xs:boolean</code>	<code>java.lang.Boolean</code>	
<code>xs:decimal</code>	<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 (historized or replicated table). However, this facet can be overwritten with a greater value in the data model.
<code>xs:dateTime</code>	<code>java.util.Date</code>	
<code>xs:time</code>	<code>java.util.Date</code>	The date portion of the returned <code>Date</code> is always set to '1970/01/01'.
<code>xs:date</code>	<code>java.util.Date</code>	The time portion of the returned <code>Date</code> is always the beginning of the day, that is, '00:00:00'.
<code>xs:anyURI</code>	<code>java.net.URI</code>	
<code>xs:Name</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>	
<code>xs:int</code> (<code>xs:decimal</code> restriction)	<code>java.lang.Integer</code>	

XML SCHEMA TYPE	JAVA CLASS	NOTES
<code>xs:integer</code> (<code>xs:decimal</code> 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](#).

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.

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®.

Extended simple types defined by EBX®

EBX® provides pre-defined simple data types:

XML SCHEMA TYPE	JAVA CLASS
<code>osd:text</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>
<code>osd:html</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>
<code>osd:email</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>
<code>osd:password</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>
<code>osd:color</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>
<code>osd:resource</code> (<code>xs:anyURI</code> restriction)	internal class
<code>osd:locale</code> (<code>xs:string</code> restriction)	<code>java.util.Locale</code>
<code>osd:dataspaceKey</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>
<code>osd:datasetName</code> (<code>xs:string</code> restriction)	<code>java.lang.String</code>

The above types are defined by the internal schema `common-1.0.xsd`. They are defined as follows:

<code>osd:text</code>	This type represents textual information. Its default user interface in EBX® consists of a dedicated editor with several lines for input and display. The use of this data type is intended for very long texts; this is why the <code>'Text'</code> search strategy is associated with this type, and cannot be modified. As a consequence, fields of type <code>osd:text</code> cannot benefit from sorting; moreover, their indexing strategy and search results differ from those of the <code>xs:string</code> type. See Default strategy for string fields for more information. <pre><xs:simpleType name="text"> <xs:restriction base="xs:string" /> </xs:simpleType></pre>
<code>osd:html</code>	This represents a character string with HTML formatting. A WYSIWYG editor is provided in EBX®. <pre><xs:simpleType name="html"> <xs:restriction base="xs:string" /> </xs:simpleType></pre>
<code>osd:email</code>	This represents an email address as specified by the RFC822 standard. <pre><xs:simpleType name="email"> <xs:restriction base="xs:string" /></pre>

	<pre></xs:simpleType></pre>
<code>osd:password</code>	This represents a hashed or encrypted password. A specific editor is provided in EBX®. <pre><xs:element name="password" type="osd:password"></pre> 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</code>. It is also possible for the default editor to use a different encryption mechanism by specifying a class that implements the interface <code>Encryption</code>. <pre><xs:element name="password" type="osd:password"> <xs:annotation> <xs:appinfo> <osd:uiBean class="com.orchestranetworks.ui.password.PasswordEditor"> <encryptionClass>package.Encryption </osd:uiBean> </xs:appinfo> </xs:annotation> </xs:element></pre> 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</code>.

	<pre><xs:element name="password" type="osd:password"> <xs:annotation> <xs:appinfo> <osd:uiBean class="com.orchestranetw <encryptionClass>package.HashCla <saltPath>../login</saltPath> </osd:uiBean> </xs:appinfo> </xs:annotation> </xs:element></pre>
<code>osd:locale</code>	This represents a geographical, political or cultural location. The locale type is translated into Java by the class <code>java.util.Locale</code>. <pre><xs:simpleType name="locale"> <xs:restriction base="xs:string"> <xs:enumeration value="ar" osd:label="Ar <xs:enumeration value="ar_AE" osd:label= <xs:enumeration value="ar_BH" osd:label= <xs:enumeration value="ar_DZ" osd:label= <xs:enumeration value="ar_EG" osd:label= <xs:enumeration value="ar_IQ" osd:label= ...</pre>

	<pre><xs:enumeration value="vi_VN" osd:label= <xs:enumeration value="zh" osd:label="Ch <xs:enumeration value="zh_CN" osd:label= <xs:enumeration value="zh_HK" osd:label= <xs:enumeration value="zh_TW" osd:label= </xs:restriction> </xs:simpleType></pre>
<code>osd:color</code>	This represents a character string with hexadecimal RGB color formatting. A color picker <code>UIComponent</code> is provided in EBX®. <pre><xs:simpleType name="color"> <xs:restriction base="xs:string" /> </xs:simpleType></pre>
<code>osd:resource</code>	This represents a resource packaged in a module. For more information, see Packaged resources. This type requires the definition of the facet FacetOResource. <pre><xs:simpleType name="resource"> <xs:restriction base="xs:anyURI" /> </xs:simpleType></pre>
<code>osd:dataspaceKey</code>	This type represents a reference to a dataspace. <pre><xs:element name="dataspaceField" type="osd:dataspaceKey"/></pre>

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:dataspaceField"/>
```

```
<xs:annotation>
```

```
<xs:appinfo>
```

```
<osd:dataspaceSet>
```

```
<include>
```

```
<pattern>a pattern</pattern>
```

```
<type>all | branch | version | none</type>
```

```
<includeDescendants>none | all</includeDescendants>
```

```
</include>
```

```
<exclude>
```

```
<pattern>a pattern</pattern>
```

```
<type>all | branch | version | none</type>
```

```
<includeDescendants>none | all</includeDescendants>
```

```
</include>
```

```
<filter osd:class="com.foo.MyData">
```

```
<param1>...</param1>
```



```
<param2>...</param2>
```

```
</filter>
```

```
</osd:dataspaceSet>
```

```
</xs:appinfo>
```

```
</xs:annotation>
```

```
</xs:element>
```

- includes

Specifies the dataspaces that can be referenced by this field. An **include** must at least be defined.

pattern : Specifies a pattern that filters dataspaces. The pattern is checked against the name of the dataspaces. This property is mandatory.

type : Specifies the type of dataspaces 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 dataspaces that match the specified pattern are included in the set. If **none** then neither children nor descendants of the dataspaces 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 <code>osd:class</code> specifies a Java bean that implements the interface <code>DataspaceSetFilter</code>. It is also possible to customize validation messages and the control policy associated with this type using the element <code>validation</code> under <code>xs:annotation/xs:appInfo/osd:dataspaceSet</code>. See Facet validation message with severity and Control policy for more information.
<code>osd:datasetName</code>	This type represents a reference to a dataset. <pre><xs:element name="dataset" type="osd:datasetName"</pre> 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 <code>osd:datasetSet</code> under <code>xs:annotation/xs:appInfo</code>: <pre><xs:element name="datasetField" type="osd:datasetName" <xs:annotation> <xs:appinfo> <osd:datasetSet> <branch>productsBranch</branch> <version>productsVersion</version> <dataspaceSelector>../dataspaceF <pattern>a pattern</pattern></pre>

```

        <filter osd:class="com.foo.MyDat

        <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.

The referred node must respect the restrictions existing for dynamic facets.

See also:

- `HomeKey.format`
- [Dynamic constraints](#)

- `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**. 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](#) and [Control policy](#) for more information.

Complex types defined by EBX®

EBX® provides pre-defined complex data types:

XML SCHEMA TYPE	DESCRIPTION
<code>osd:UDA</code>	User Defined Attribute: 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.
<code>osd:UDACatalog</code>	Catalog of User Defined Attributes: This type consists of a table in which attributes can be specified. This catalog is used by all <code>osd:UDA</code> elements declared in the same data model.

<code>osd:UDA</code>	A User Defined Attribute (UDA) supports both the <code>minOccurs</code> and <code>maxOccurs</code> attributes, as well as the attribute <code>osd:UDACatalogPath</code>, which specifies the path of the corresponding catalog. <pre><xs:element name="firstUDA" type="osd:UDA" minOccurs="1" maxOccurs="unbounded" osd:UDACatalogPath="//UDACatalog" /> <xs:element name="secondUDA" type="osd:UDA" minOccurs="1" maxOccurs="1" osd:UDACatalogPath="/root/userCatalog" /> <xs:element name="thirdUDA" type="osd:UDA" minOccurs="1" maxOccurs="1" osd:UDACatalogPath="/root/userCatalog" /></pre>

	<pre> osd:UDACatalogPath="//userCatalog" /></pre> In the manager, when working with a UDA, the editor will adapt itself to the type of the selected attribute.
<code>osd:UDACatalog</code>	Internally, a catalog is represented as a table. The parameters <code>minOccurs</code> and <code>maxOccurs</code> must be specified. Several catalogs can be defined in the same data model. <pre><xs:element name="insuranceCatalog" type="osd:UDACatalog"> <xs:annotation> <xs:documentation xml:lang="en-US">Insurance Catalog</xs:documentation> <xs:documentation xml:lang="fr-FR">Catalogue d'assurance</xs:documentation> </xs:annotation> </xs:element> <xs:element name="userCatalog" type="osd:UDACatalog"> <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></pre> Only the following types are available for creating new attributes: • <code>xs:string</code>

- 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.

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: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](#).

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` 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.

Tables and relationships

- [Tables](#)
 - [Overview](#)
 - [Declaration](#)
 - [Common properties](#)
 - [Example](#)
 - [Properties related to dataset inheritance](#)
 - [Using toolbars](#)
 - [Example](#)
 - [Record metadata](#)
- [Foreign keys](#)
 - [Declaration](#)
 - [Example](#)
- [Associations](#)
 - [Overview](#)
 - [Declaration](#)
 - [User interface integration](#)
 - [Validation](#)
 - [Data services integration](#)
 - [Examples](#)
- [Linked fields](#)
 - [Overview](#)

- Declaration

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](#). Beyond relational capabilities, some features that tables provide are:

- filters and searches;
- sorting, views and hierarchies;
- identity constraints: primary keys, [foreign keys](#) and [uniqueness constraints](#);
- 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](#));
- performance and memory optimization.

See also:

- [Entity mappings](#)
- [Foreign keys](#)
- [Associations](#)
- [Linked fields](#)
- [Working with existing datasets](#)
- [Simple tabular views](#)
- [Hierarchical views](#)
- [History](#)

Declaration

A table element, which is an element with *maxOccurs* > 1, is declared by adding the following annotation:

```
<xs:annotation>

  <xs:appinfo>

    <osd:table>

      <primaryKeys>/pathToField1 /pathToField...n</primaryKeys>

    </osd:table>

  </xs:appinfo>

</xs:annotation>
```

Note:

A default entity name is associated to the table if it does not have one explicitly set in the data model. The table's name is used as its default entity name. Entity names allow to refer to tables in SQL requests and Data Service operations using unique names instead of their paths by defining suffixes for WSDL operations and aliases for SQL requests. See [Entity mappings](#) for more informations.

Common properties

ELEMENT	DESCRIPTION	REQUIRED
<code>primaryKeys</code>	Specifies the primary key fields of the table. 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	Yes.

ELEMENT	DESCRIPTION	REQUIRED
	are multiple fields in the primary key, the list is delimited by whitespace. Note: Whitespaces in primary keys of type <code>xs:string</code> are handled differently. See Whitespace handling for primary keys of type string.	
<code>defaultLabel</code>	Defines the end-user display of records. Multiple variants can be specified: - A static non-localized expression is defined using the <code>defaultLabel</code> element, for example: <pre><defaultLabel>Product : \${./productCode}</ defaultLabel></pre> - Static localized expressions are specified using the <code>defaultLabel</code> element with the attribute <code>xml:lang</code>, for example: <pre><defaultLabel xml:lang="fr - FR">Produit : \${./</pre>	No.

ELEMENT	DESCRIPTION	REQUIRED
	<pre>productCode}</ defaultLabel> <defaultLabel xml:lang="en- US">Product: \${./ productCode}</ defaultLabel></pre> A JavaBean that implements the interface <code>UILabelRenderer</code> and/or the interface <code>UILabelRendererForHierarchy</code>. The JavaBean is specified by means of the attribute <code>osd:class</code>, for example: <pre><defaultLabel osd:class="com.wombat.MyLabel"></ defaultLabel></pre> Attention: Since the end-user display should be sortable, only fields that use sortable search strategies are allowed in static expressions. Note: The priority of the tags when displaying the user interface is the following:	

ELEMENT	DESCRIPTION	REQUIRED
	1. <code>defaultLabel</code> tags with a JavaBean (but it is not allowed to define several renderers of the same type); 2. <code>defaultLabel</code> tags with a static localized expression using the <code>xml:lang</code> attribute; 3. <code>defaultLabel</code> tags with a static non-localized expression. Attention: 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.	
<code>recordForm</code>	Defines a specific component for customizing the record form in a dataset. This component is defined using a JavaBean that extends <code>UIForm</code> or implements <code>UserServiceRecordFormFactory</code>. The JavaBean is specified by means of the attribute <code>osd:class</code>, for example:	No.

ELEMENT	DESCRIPTION	REQUIRED
	<pre><recordForm osd:class="com.wombat.MyRecordForm"/ ></pre>	

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: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
<code>onDelete-</code> <code>deleteOccultingChildren</code>	Specifies whether, upon record deletion, child records in occulting mode are also to be deleted. Valid values are: <code>never</code> or <code>always</code> .	No, default is <code>never</code> .

ELEMENT	DESCRIPTION	REQUIRED
<code>mayCreateRoot</code>	Specifies whether root record creation is allowed. The expression must follow the syntax below. See definition modes .	No, default is <code>always</code> .
<code>mayCreateOverwriting</code>	Specifies whether records are allowed to be overwritten in child datasets. The expression must follow the syntax below. See definition modes .	No, default is <code>always</code> .
<code>mayCreateOcculting</code>	Specifies whether records are allowed to be occulted in child datasets. The expression must follow the syntax below. See definition modes .	No, default is <code>always</code> .
<code>mayDuplicate</code>	Specifies whether record duplication is allowed. The expression must follow the syntax below.	No, default is <code>always</code> .
<code>mayDelete</code>	Specifies whether record deletion is allowed. The expression must follow the syntax below.	No, default is <code>always</code> .

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 ::=
```

"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](#)

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](#)

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.

Record metadata

EBX® automatically adds record metadata to a table's structure. The metadata nodes enable retrieval of the following information about a record:

- creator
- creation time
- last modifier
- time of last modification
- Unique identifier of a record (UUID). This metadata being technical it cannot be used in a custom table view, in the structured search tools and quick search.

An example use for record metadata might be to display or fetch all the records created by a specific user since a specific date. This is possible:

- For business users in the user interface by:
 - Using the quick search feature to locate records by using the metadata fields as part of the search criteria.
 - Defining table views that include the desired metadata fields.
- Via the API by querying with XPath or SQL.

Attention:

Please take note of the following points regarding metadata:

- Metadata are defined under a group with the reserved name *ebx-metadata*. As a consequence, the metadata won't be available in a table if it contains a group, directly under its root, with the same name. A warning displays when compiling the data model if a table defines a group with this reserved name.
- Metadata are read-only.
- Metadata are not included when exporting the content of a table to an XML or CSV file.
- By default, metadata fields are not including when introspecting a data model using the API. You can explicitly included them using the dedicated [API](#).

See also:

- [Metadata fields in SQL](#)
- [Technical data](#)
- `SchemaMetadataConstants`
- `Constants.Data`

Foreign keys

Declaration

A reference to a [table](#) 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](#).

This extended facet is also interpreted as an enumeration whose values refer to the records in the target table.

ELEMENT	DESCRIPTION	REQUIRED
<code>tablePath</code>	XPath expression that specifies the target table.	Yes.
<code>container</code>	Reference of the dataset that contains the target table.	Only if the dataspace element is defined. Otherwise, default is the current dataset.
<code>branch</code>	Reference of the dataspace that contains the container dataset.	No, default is the current dataspace or snapshot.
<code>display</code>	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:	No, if the <code>display</code> property is not specified, the table's record rendering is used.

ELEMENT	DESCRIPTION	REQUIRED
	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> <display> <pattern>Product : \${./productCode}</ pattern> <pattern xml:lang="fr- FR">Produit : \${./ productCode}</ pattern> <pattern xml:lang="en- US">Product: \${./ productCode}</ pattern> </display> </pre> Attention: Since the display pattern should be sortable, only fields that use sortable	

ELEMENT	DESCRIPTION	REQUIRED
	search strategies are allowed. A JavaBean that implements the interface <code>TableRefDisplay</code>. It is specified using the attribute <code>osd:class</code>. For example: <pre><display osd:class="com.wombat.MyLabel"></display></pre> Attention: Quick search and sort operations in the user interface will use the raw value instead of the label of the records if a JavaBean is defined or if the target table defines a programmatic label for its records. A static expression must be defined to use in these operations the label of the records. It is not possible to define both variants on the same foreign key element. Attention: Access rights defined on associated datasets are not applied when displaying record	

ELEMENT	DESCRIPTION	REQUIRED
	labels. Fields that are usually hidden due to access rights restrictions will be displayed in labels.	
<code>filter</code>	Specifies an additional constraint that filters the records of the target table. Two types of filters are available: An XPath filter is an XPath predicate in the target table context. It is specified using the <code>predicate</code> element. For example: <pre><filter><predicate>type = \${../refType}</predicate></filter></pre> 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. A specific severity level can be defined in a nested <code>severity</code>	No.

ELEMENT	DESCRIPTION	REQUIRED
	element. The default severity is 'error'. 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 <code>message</code> element with no <code>xml:lang</code> attribute. 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.	

ELEMENT	DESCRIPTION	REQUIRED
	A programmatic filter is a JavaBean that implements the interface <code>TableRefFilter</code>. It is specified using the attribute <code>osd:class</code>. For example: <pre><filter osd:class="com.wombat.MyFilter"></ filter></pre> Additional validation messages can be specified during the setup of the programmatic filter. 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	

ELEMENT	DESCRIPTION	REQUIRED
	the method <code>TableRefFilterContext.setFilterForSearch</code>. Note: The attributes <code>osd:class</code> and the property <code>predicate</code> cannot be set simultaneously. The validation search XPath functions are forbidden on a <code>tableRef</code> filter. See also: • JavaBean specifications • <code>JavaBeanVersion</code>	
<code>validation</code>	Specifies localized validation messages for the <code>osd:tableRef</code> and error management policy. A specific severity level can be defined in a nested <code>severity</code> element. The default severity is 'error'. 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	No.

ELEMENT	DESCRIPTION	REQUIRED
	key filter. In this case, updates are not rejected, and a validation error will be reported. 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 <code>message</code> element with no <code>xml:lang</code> attribute.	

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](#)
- [Primary key syntax](#)
- [Extraction of foreign keys \(XPath predicate syntax\)](#)
- [Associations](#)
- [View for advanced selection](#)
- `SchemaNode.getFacetOnTableReference`
- `SchemaFacetTableRef`

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](#).
- *'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](#).
- 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`
- `SchemaNode.isAssociationNode`
- `AssociationLink`

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
<code>tableRefInverse</code>	Defines the properties of an association that is the inverse relationship of a <i>foreign key</i>. 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. The element <code>fieldToSource</code> can be defined on a foreign key list (<code>maxOccurs</code> > 1), in that case, the list should be declared as unique with a blocking uniqueness constraint (<code>onInsertUpdateOrDelete</code>). See also: • Blocking and non-blocking constraints • Uniqueness constraints	Yes if the association is the inverse relationship of a <i>foreign key</i> , otherwise no.
<code>linkTable</code>	Defines the properties of an association over a link table.	Yes if the association is over a link table, otherwise no.

ELEMENT	DESCRIPTION	REQUIRED
	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. Important: 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. 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. 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.	

ELEMENT	DESCRIPTION	REQUIRED
<code>xpathLink</code>	Defines the properties of an association that is based on an <i>XPath predicate</i>. The <code>predicate</code> element specifies the criteria of the association, relative to the current node. Examples: <code>/root/Products[catalog_ref =\${../catalogId}]</code> or <code>//Products[catalog_ref =\${../catalogId}]</code> or <code>../Products[catalog_ref =\${../catalogId}]</code>. 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. 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>\${<relative-path>}</code> where <code>relative-path</code> is a path that	Yes if the association is based on an <i>XPath predicate</i> , otherwise no.

ELEMENT	DESCRIPTION	REQUIRED
	identifies the element relative to the association node. See EBX® XPath supported syntax. Note: The validation search XPath functions are forbidden on an XPath link.	
<code>filter</code>	Defines an XPath predicate to filter associated objects using the <code>predicate</code> element. For example: <pre><filter><predicate>type = \${../refType}</predicate></filter></pre> 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. Error message on creation: in the user interface, the record creation is blocked when a user submits a new associated record that does not comply with the filter.	No.

ELEMENT	DESCRIPTION	REQUIRED
	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 Examples for more information on this property. Note: The validation search XPath functions are forbidden for association filter.	
<code>xpathFilter</code>	Note: Deprecated. This property has been replaced by the property <code>filter</code>. Defines an XPath predicate to filter associated objects.	No.
<code>recordForm</code>	Defines a specific component for customizing the form of an associated record. This component is defined using a JavaBean	No.

ELEMENT	DESCRIPTION	REQUIRED
	that implements <code>UserServiceAssociationRecordFormFactory</code>. The JavaBean is specified by means of the attribute <code>osd:class</code>, for example: <pre><recordForm osd:class="com.wombat.MyRecordFormFactory"/ ></pre>	

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</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.
- The consistency of an association to another dataset is checked in the operational phase. If an association is not consistent in the context of a specific dataset then:
 - It is considered as hidden from a permission perspective. That is, the association is hidden in the user interface and in data services as well.
 - It is excluded from the validation of the container dataset, table or record.
 - An exception is raised in the existing API that compute the result of the association.

See `AssociationLink.getAssociationResult` and
`AssociationLink.getTargetTable`

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
<code>tabularViewTop</code>	Defines the toolbar to use in the default table view of this association.	No.
<code>tabularViewRow</code>	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](#)

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
<code>column</code>	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 <code>column</code> elements can be defined to specify the fields that are to be displayed.	No.
<code>sort</code>	Define a field that can be used to sort associated objects. Several <code>sort</code> elements can be defined to specify the fields that can be used to sort associated objects. The element <code>nodePath</code> defines the path of the field that can be used to sort associated objects. The element <code>isAscending</code> specifies whether the sort	No.

ELEMENT	DESCRIPTION	REQUIRED
	order is ascending (true) or descending (false).	

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	No, by default the minimum is not restricted.

ELEMENT	DESCRIPTION	REQUIRED
	element <code>value</code> and must be a positive integer.	
<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>	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. 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.	No.
<code>validation</code>	A validation message can be defined under the elements	No.

ELEMENT	DESCRIPTION	REQUIRED
	<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> . If the severity is not specified then, by default, the severity <code>error</code> is used. 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.	

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
- Association field

Examples

For example, the product catalog data model defined [previously](#) 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>

```

Linked fields

Overview

Linked fields provide the ability to simulate a multi-table view by aggregating the fields of a main table with some fields of another table over an existing relationship. It allows an easy model-driven integration of some fields from records that are referred by other ones using relationships.

See also:

- `SchemaLinkedField`

Following relationships can be traversed:

- [Foreign key constraints](#).
- [Associations 'By foreign key'](#).

Important: these relationships must be "single valued relationships". That is, a record of the table that defines the relationship must refer to only one record in the target table of the relationship.

In details:

- A foreign key is considered as a single valued relationship if the foreign key field defines a *maxOccurs* equal to 1.

- An association by foreign key is considered as a single valued relationship if the referred foreign key field is the unique primary key field of the target table or if an uniqueness constraint is defined on this foreign key field.

Declaration

A linked field, which is an element with *minOccurs* = 0 and *maxOccurs* = 0, is declared by adding the following annotation:

```
<xs:element name="catalog_ref_label" minOccurs="0" maxOccurs="0"
  type="xs:string">

  <xs:annotation>

    <xs:appinfo>

      <osd:function linkedField="../../catalog_ref/catalogLabel"/>

    </xs:appinfo>

  </xs:annotation>

</xs:element>
```

Attribute *linkedField* defined in element *osd:function* defines a path composed of steps which refer to a single-valued relationship in the container table and of steps to a field in the target table of the relationship.

Important:

- If the path is absolute (starts with a "/") then it will be resolved from the container table node.
- If the path is relative then it will be resolved from the current field.

In the context of the product catalog data model defined [previously](#), the field *catalog_ref_label* is added in the *Products* table. The first step *catalog_ref* refers to the foreign key that indicates that a product belongs to a catalog. The last step *catalogLabel* refers to the label of the referred catalog. This field belongs to the *Catalog* table that is the target table of the foreign key constraint.

Restrictions:

- Currently, linked fields only allow to target one level of single valued relationships. As a consequence:
 - The single-valued relationship cannot be defined by another linked field.
 - A linked field cannot target another linked field.
- A linked field must define properties *minOccurs* and *maxOccurs*. equal to 0.
- A linked field must define a data type that is compatible with the one of the target field.
- A linked field cannot be a part of the primary key of the container table.
- A linked field must target a terminal field.
- A linked field cannot target an aggregated list or a field under an aggregated list.
- A linked field cannot refer to a single-valued relationship that is also a linked field.
- Targeting a field which is a static enumeration with localized labels (defines *xs:enumeration* in the data model) has some limitations. That is, only the raw value of the linked field will be available for display and search, instead of the enumeration label of the target field. It is still possible to define the same static enumerations on both the linked field and its target to benefit of the labelling features for display and search.
- If a linked field and its target define both a foreign key constraint then these constraints must define exactly the same display pattern (property *osd:tableRef/display/pattern* in the data model). Otherwise the display pattern defined locally on the linked field will be ignored during a search operation. That is, the search operation works only the display pattern defined by the targeted foreign key field.

Constraints

Facets allow you to define data constraints in your data models. TIBCO EBX® supports XML Schema constraining facets and provides extended and programmatic facets for advanced data controls.

- [XML Schema supported facets](#)
 - [Uniqueness constraint](#)
- [Extended facets](#)
 - [Foreign keys](#)
 - [Dynamic constraints](#)
 - [FacetOResource constraint](#)

- Excluding values
- Enumeration constraint defined using another node
- Programmatic facets
 - Programmatic constraints
 - Programmatic enumeration constraints
 - Constraint on 'null' values
 - Constraints on table
- Control policy
 - Blocking and non-blocking constraints
 - Check 'null' input
 - EBX® whitespace management for data types
 - Empty string management
 - Validation message threshold

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](#).
- **4** - `osd:dataspaceKey`, `osd:datasetName` and `osd:color` types do not support facets.

Only [Programmatic constraints](#) are supported on these types.

	LENGTH	MINLENGTH	MAX LENGTH	PATTERN	ENUMERATION	WHITE SPACE
<code>xs:string</code>	X	X	X	X	X	1

	LENGTH	MINLENGTH	MAX LENGTH	PATTERN	ENUMERATION	WHITE SPACE
xs:boolean				X		1
xs:decimal				X	X	1
xs:dateTime				X	X	1
xs:time				X	X	1
xs:date				X	X	1
xs:anyURI	X	X	X	X	X	1
xs:Name	X	X	X	X	X	1
xs:integer				X	X	1
osd:resource 3						1
osd:dataspaceKey 4						1
osd:datasetName 4						1
osd:color 4						1

	FRACTION DIGITS	TOTAL DIGITS	MAX INCLUSIVE	MAX EXCLUSIVE	MIN INCLUSIVE	MIN EXCLUSIVE
xs:string			2	2	2	2
xs:boolean						
xs:decimal	X	X	X	X	X	X
xs:dateTime			X	X	X	X
xs:time			X	X	X	X

	FRACTION DIGITS	TOTAL DIGITS	MAX INCLUSIVE	MAX EXCLUSIVE	MIN INCLUSIVE	MIN EXCLUSIVE
<code>xs:date</code>			X	X	X	X
<code>xs:anyURI</code>						
<code>xs:Name</code>			2	2	2	2
<code>xs:integer</code>	X	X	X	X	X	X
<code>osd:resource</code> 3						
<code>osd:dataspaceKey</code> 4						
<code>osd:datasetName</code> 4						
<code>osd:color</code> 4						

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
<code>name</code> attribute	Identifies the constraint in the data model.	Yes
<code>xs:selector</code> 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
<code>xs:field</code> element	Indicates the field in the context whose values must be unique, using a restricted XPath expression. It is possible to indicate that a set of values must be unique by defining multiple <code>xs:field</code> 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.

The uniqueness constraint can also be applied on a simple aggregated list: in this case, each value of the list has to be unique in the scope of the list and not in the scope of the table.

Example:

In the example below, a uniqueness constraint is defined on the 'title' table, for the target field 'printedEditions'. This means that an edition can appear only once in the list.

```
<xs:element name="title">
  ...
  <xs:complexType>
    <xs:sequence>
      ...
      <xs:element name="printedEditions" type="xs:string"
minOccurs="0" maxOccurs="5"/>
      ...
    </xs:sequence>
  </xs:complexType>
  <xs:unique name="uniquePrintedEditions">
    <xs:annotation>
```

```

        <xs:appinfo>

            <osd:validation>

                <severity>error</severity>

                <message xml:lang="en-US">An edition must be
referenced only once by this title</message>

                <message xml:lang="fr-FR">Une édition ne peut être
référéncée qu'une seule fois par ce livre</message>

            </osd:validation>

        </xs:appinfo>

    </xs:annotation>

    <xs:selector xpath="." />

    <xs:field xpath="printedEditions"/>

</xs:unique>

</xs:element>

```

Limitations:

1. The target of the `xs:field` element must be in a table.
2. The uniqueness constraint does not apply to computed fields.
3. The uniqueness constraint cannot be applied on multiple fields that contains an aggregated list.
4. The uniqueness constraint cannot be applied on embedded lists.

See also:

- [Uniqueness constraint in the Java API](#)

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](#) 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>
```



```

        </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](#). It has the following attributes:

`moduleName`

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><dependencies></code> in the file <code>module.xml</code> .
<code>resourceType</code>	Represents the resource type that is one of the following values: 'Image', 'JavaScript', 'Style sheet', 'HTML'.
<code>relativePath</code>	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](#).

Excluding values

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:element>

```

```

        </xs:appinfo>

    </xs:annotation>

    <xs:simpleType type="xs:string" />

</xs:element>

```

Enumeration constraint defined using another node

Attention:

This kind of constraint is obsolete. You should use [foreign key constraint](#). It has limitations, in particular Quick search and sort operations in the user interface will use the raw value of the field instead of the labels defined through the constraint.

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:sequence>

  </xs:complexType>

</xs:element>
```

```

        </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>

```

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`.

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](#)
- `JavaBeanVersion`

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`.

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>

```

Attention:

Quick search and sort operations in the user interface will use the raw value of the field instead of its label. Consider whether this enumeration can be replaced by a [foreign key constraint](#) to a table that defines the same set of values.

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`.
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 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>
```

See also:

- [JavaBean specifications](#)
- `JavaBeanVersion`

Constraints on table

A constraint on table is defined by a Java class that implements the interface `ConstraintOnTable`. 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>

      <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 [Data validation](#).

See also:

- [JavaBean specifications](#)
- `JavaBeanVersion`

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:

<code>onInsertUpdateOrDelete</code>	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. 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.
<code>onUserSubmit-checkModifiedValues</code>	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. 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.
<code>never</code>	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.

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 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](#)
- [Foreign keys](#)

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>
```

```

        </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.setBlocksCommitToNever` and `ConstraintContextOnTable.setBlocksCommitToNever` 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>

```

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"`. This property is ignored if defined on an aggregated list (`maxOccurs` > 1).

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'
- Whitespace management
- Empty string management

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*:

preserve	No normalization is performed, the value is unchanged.
replace	All occurrences of <code>#x9</code> (tab), <code>#xA</code> (line feed) and <code>#xD</code> (carriage return) are replaced with <code>#x20</code> (space).
collapse	After the processing according to the replace procedure, contiguous sequences of <code>#x20</code> are then collapsed to a single <code>#x20</code> , and any leading or trailing <code>#x20</code> s 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](#)

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.

Validation message threshold

It is possible to specify at the data model level the maximum number of validation messages allowed per constraint when performing a validation.

Example:

```
<xs:schema ...>

    ...

    <xs:annotation>

        <xs:appinfo>

            <osd:validation>

                <validationMessageThreshold>250</validationMessageThreshold>

            </osd:validation>

        </xs:appinfo>

    </xs:annotation>

    ...

</xs:schema>
```

The threshold is considered for each constraint defined in a data model and in each dataset validation report. When the threshold is reached by a constraint, the validation of the constraint is stopped and an error message indicating that the threshold has been reached is added to the validation report.

The validation message threshold is set by default to `1000` if it is not defined in the data model. It is not allowed to set an unlimited number of validation messages. Also, the specified validation message threshold must be greater or equal than 100.

See also:

- `ValidationReport.hasConstraintsWithTooManyMessages`

Triggers and functions

EBX® data model allows to define triggers and computed fields. It also provides auto-incremented fields

- [Computed values](#)
 - [Disabling validation](#)
- [Triggers](#)
 - [Trigger on dataset](#)
 - [Trigger on table](#)
- [Auto-incremented values](#)

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 *computation rules*.

A computation rule 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 `ValueFunction`
- 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>

```

Disabling validation

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">

```

```

        ...

        </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.

Trigger on dataset

For dataset triggers, a Java class that extends the abstract class `InstanceTrigger` 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>
            </osd:trigger>
        </xs:appinfo>
    </xs:annotation>
</xs:element>

```



```

        <param...n>...</param...n>

    </osd:trigger>

</xs:appinfo>

</xs:annotation>

...

</xs:element>

```

Trigger on table

For the definition of table record triggers, a Java class that extends the abstract class `TableTrigger` 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](#).

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.

- [Label and description](#)
 - [Dynamic labels and descriptions](#)
- [Enumeration labels](#)
- [Mandatory error message \(osd:mandatoryErrorMessage\)](#)
- [Conversion error message](#)
- [Facet validation message with severity](#)
 - [XML Schema facet \(osd:validation\)](#)
 - [XML Schema enumeration facet \(osd:enumerationValidation\)](#)
 - [EBX® facet \(osd:validation\)](#)

Label and description

A label and a description can be added to each node in a data model. In EBX®, each 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 <code><osd:label></code> and <code><osd:description></code> 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`

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>

```

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.

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>

```

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>

```

Additional properties

- [Default values](#)
- [Access properties](#)
- [Information](#)
- [Default view](#)
 - [Hiding a field or a table in the default view](#)
 - [Hiding groups and fields in views](#)
 - [Hiding a field in structured search tools](#)

- Hiding a field in Data Services
- Defining a view for the combo box selector of a foreign key
- Customizing a default widget
- Customizing REST data services
- Comparison mode
- Apply last modifications policy
- Categories
 - Filters on data
 - Predefined categories
 - Restriction

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" />
```

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`

ACCESS MODE	BEHAVIOR
<code>RW</code>	<i>Adaptation terminal node</i> : value can be read and written in EBX®.
<code>R-</code>	<i>Adaptation terminal node</i> : value can only be read in EBX®.
<code>CC</code>	<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 access attribute.

ACCESS MODE	BEHAVIOR
--	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. The <code>root</code> node of a data model must specify this access mode.
Default	If the access attribute is not defined: • If the node is a computed value, it is considered to be <code>R-</code> • If the node is a simple type and its value is not computed, it is considered to be <code>RW</code> • If the node is an aggregated list, it is then a terminal value and is considered to be <code>RW</code> • Otherwise, it is not an adaptation terminal node and it does not define anything about its child nodes.

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"/>
```

Information

The element `osd:information` allows specifying additional information.

See also:

- `SchemaNode.getInformation`

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>
```

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 structured search tools

To specify whether or not to hide an element in structured 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.

In all cases, the field will remain searchable in the quick search tool. A field can be excluded from all search tools, including the quick search, by defining a specific search strategy.

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: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](#).

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](#) for more information.

Customizing a default widget

A widget can be defined using the data model assistant. See [Default view > Widget](#) for more information.

Customizing REST data services

Default view configuration is managed in REST data services through the [session channel](#).

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:

default	Element is visible during comparisons of its data.
ignored	No changes will be detected when comparing two versions of the content in records or datasets. 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. 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.

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"/>
```

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:

default	Last modifications can be applied to this element.
ignored	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"/>
```

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 > ****** 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'.

Data services

This chapter details how WSDL operations' names related to a table are defined and managed by TIBCO EBX®.

- [Definition](#)
- [Configuration](#)
 - [Embedded data model](#)
 - [Packaged data model](#)

- [Publication](#)
- [WSDL and table operations](#)
 - [WSDL Generator](#)
 - [SOAP operations](#)
- [Limitations](#)

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 entity names. Entity names allow to use unique names instead of the last element of tables' paths. Entity names override the suffix operations' names.

See also:

- [Entity mappings using the Data Model Assistant](#)

Configuration

Embedded data model

Entity names used as 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 entity names can be defined in the data model definition, under the 'Configuration > 'Entity mappings' table and managed by EBX®.

Packaged data model

Entity names used as 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.

Publication

Entity 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.

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](#)

Limitations

WSDL operations' names are not available with `external` data models.

Toolbars

This chapter details how toolbars are defined and managed by TIBCO EBX®.

- [Definition](#)
- [Using toolbars](#)

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](#)
- [Using toolbars in data models](#)
- [Toolbar API](#)

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 the top of a tile view.
- On each card of a tile 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](#)
- [Associations](#)

Custom forms

This chapter details how custom forms are defined and managed by TIBCO EBX®.

Related concepts:

- [Interface customization](#)
- [Access](#)
- [Forms and components](#)
- [The editor](#)
- [Blocks](#)
 - [Arithmetic operator](#)
 - [Array](#)
 - [Array item](#)
 - [Boolean operator](#)


- Check box
- Check box group
- Children
- Choose
- Close button
- Color picker
- Combo box
- Comparator
- Comparison
- Concatenate arrays
- Constant
- Content title
- Creating record?
- Custom action bar
- Date time input
- Declare
- Default action bar
- Default toolbar
- Duplicating record?
- Expand/Collapse
- Flow
- For each
- Form
- Form definition
- Form group
- Form row
- Form parameter
- From...get
- Function
- Function call

- Function call argument
- Function parameter
- Grid
- Grid element
- If
- Input parameter
- In role?
- Label
- Localized
- Message
- Next button
- Node
- Nomenclature item
- Not
- Object
- Otherwise
- Paragraph title
- Password
- Previous button
- Property
- Radio button
- Radio button group
- Revert button
- Save button
- Save and close button
- Tab
- Tab container
- Text input
- Toolbar named
- Variable


- When
- Widget
- With

Access

To access it, go to: Data Models > Extensions > Custom forms > Forms




Click on the button to create a form.



Actions ▾ No records ▾

Name	Target table
No records found.	


Name your new form and indicate the table on which it will be available.



Name  AlbumForm

Target table  /root/artist ▾

Pressing the "Save" button will redirect you to the layout designer.



Actions ▾

- 1 - Specification
- 2 - Layout
- 3 - Component
- Context
- Operator
- Data
- Custom component
- [Templates] Specification
- [Templates] Layout
- [Templates] Component
- [Templates] Operator
- [Templates] Data

Form AlbumForm with

declare record as form para...

return

Form definition [toolbar] Default toolbar

Form

- var record/id
- var record/name
- var record/parent
- var record/child_order
- var record/description

Forms and components

Both forms and components can be created in the 'Custom forms' data model extension.

- **A form describes the layout of a record.** It can access contextual information such as the record or the input parameters. When creating a form, the user is asked to provide its target table. Once created, **it can be declared as the default form of the table** either in the 'Default form' table of the extension, or in the 'Extensions' tab of the target table. Forms that are not the default for their table will not be used.
- **A component is a reusable fragment** that can be shared between forms and other components. Unlike forms, components don't have access to contextual information. If such information is needed, it must be provided explicitly by its caller.


The editor

A form or component is just an imbrication of blocks. Blocks can represent a concrete graphical element or some piece of logic, allowing to have different layouts based on conditions like permissions, language, etc.

The workspace contains the description of the form, inside a predefined root block. Any block that is not connected to this root is grayed out, to mark it as inactive. **To move a block**, it must be dragged and dropped. Dragging a block also drags the blocks connected below it. If only one block has to be moved, hold the control key before clicking on it. **Right-clicking** on a block in the workspace also shows a list of options such as expand/collapse, comment, help, etc.

The toolbox on the left displays a list of categories. By clicking on a category, the blocks it contains are displayed. Some categories have a related 'Template' section. This section provides some combinations of the blocks of the section and can be considered as useful shortcuts or samples.

On the bottom right of the screen are the 'Preview' and 'Configure preview' buttons. These can be used to see what the form will look like.




Blocks

Here is the list of all the blocks that appear in the toolbox.


Arithmetic operator

Operations on two integers.



Array

A list of items.



See also:

- [Array item](#)

Array item


Adapter to make expression blocks into arrays.

See also:

- [Array](#)


Boolean operator

Operations on two booleans.



Check box

Displays the checkbox widget.




Parameters

INPUT	DESCRIPTION
<code>node</code>	The node for which to display the widget.
<code>readonly?</code>	Indicates if the widget should be in readonly mode.
<code>label enabled?</code>	If set, indicates if the item label is to be added next to the widget.
<code>index</code>	The index for this enumeration item.
<code>specific nomenclature</code>	If set, overrides the model-driven nomenclature.

Check box group

Displays the checkbox group widget.



Parameters

INPUT	DESCRIPTION
<code>node</code>	The node for which to display the widget.
<code>readonly?</code>	Indicates if the widget should be in readonly mode.

INPUT	DESCRIPTION
columns number	If set, defines the number of columns to use for the layout of the checkboxes.
specific nomenclature	If set, overrides the model-driven nomenclature.

Children

Returns the list of the children of the given node.




Parameters

INPUT	DESCRIPTION
node	The complex node from which to extract children.

Choose

A block returning the content of one of its inner 'when'/'otherwise' blocks.



Parameters

INPUT	DESCRIPTION
content	A list of 'when' blocks and optionally a final 'otherwise' block.

See also:

- [When](#)
- [Otherwise](#)

Close button

Standard 'Close' button.

Close button [label] [default?] false

Parameters

INPUT	DESCRIPTION
label	If set, overrides the default label.
default?	Indicates if this is the action to trigger when pressing 'Enter'. Only one button should be the default one.

Color picker

Displays the color picker widget.

node

[base] var some/node

[mandatory?] true

[label] ...

[description] ...

[widget] Color picker [readonly?] false


Color picker [node] var some/node [readonly?] false

Parameters

INPUT	DESCRIPTION
<code>node</code>	The node for which to display the widget.
<code>readonly?</code>	Indicates if the widget should be in readonly mode.

Combo box

Displays the combo box widget.




Parameters

INPUT	DESCRIPTION
<code>node</code>	The node for which to display the widget.
<code>readonly?</code>	Indicates if the widget should be in readonly mode.
<code>create button displayable?</code>	If set, defines if the create button should be displayed when the underlying node is a foreign key.
<code>preview button displayable?</code>	If set, defines if the preview button should be displayed when the underlying node is a foreign key.
<code>number of entries displayed on a page</code>	If set, defines the number of entries on each page of the drop-down list.
<code>width of the combo box</code>	If set, defines the width of the combo box.
<code>advanced selector available?</code>	If set, defines if the advanced selector should be displayed when the underlying node is a foreign key.
<code>view name to use with the advanced selector</code>	If set, defines the name of the published view that will be used in the combo-box selection of the associated foreign key field.
<code>specific nomenclature</code>	If set, overrides the model-driven nomenclature.

Comparator

Compares two integers.




Comparison

Compares strings, numbers or booleans.



Concatenate arrays

Returns the concatenation of the given arrays.



Parameters

INPUT	DESCRIPTION
arrays	A list of arrays.

Constant

A constant value. Depending on the context, this value may be interpreted as text, boolean or number.



Parameters

INPUT	DESCRIPTION
value	Text, boolean or number.

Content title

Display the workspace content title.

Content title [label] Title to display [horizontal alignment] start

Parameters

INPUT	DESCRIPTION
label	The text to display.
horizontal alignment	The horizontal alignment of the element. If set, accepted values are: 'start', 'end', 'center'.


Creating record?

Returns a boolean indicating if the form is displayed in the context of a record creation or duplication.

creating record?

Custom action bar

A custom action bar.




Parameters

INPUT	DESCRIPTION
start	A list of 'Button's to display on the left of the action bar.
middle	A list of 'Button's to display on the center of the action bar.

INPUT	DESCRIPTION
end	A list of 'Button's to display on the right of the action bar.

Date time input

Displays the date/time input.



Parameters

INPUT	DESCRIPTION
node	The node for which to display the widget.
readonly?	Indicates if the widget should be in readonly mode.
date picker button displayable?	Indicates if the button opening the date picker is displayed (read/write mode only).
now button displayed?	Indicates if the button setting the date to the current time is displayed (read/write mode only).

Declare

The declaration of a variable, in a 'with' block.

```
declare isAllowed as in role? [administrator]
```

Parameters

INPUT	DESCRIPTION
name	The name of the variable. Must be unique for a given 'with' block.
content	The value of the variable, which will be returned by 'var' blocks referencing this.

See also:

- [With](#)

Default action bar

The default action bar.

```
Default action bar
```

Default toolbar

The model-driven toolbar.

```
Default toolbar
```


Duplicating record?

Returns a boolean indicating if the form is displayed in the context of a record duplication.

```
duplicating record?
```

Expand/Collapse

Displays its content inside an expand/collapse block.




Parameters

INPUT	DESCRIPTION
label	Label displayed next to the arrow.
collapsed?	Indicates if the group is initially collapsed or not.
content	A list of components to be collapsible.

Flow

Displays its content in a fluid manner. If the elements don't fit in one row or column, they will wrap to start a new one.



Parameters

INPUT	DESCRIPTION
direction	'horizontal' or 'vertical', indicates in which direction to queue its content.
item spacing	The space between elements, in pixels.

Parameters


INPUT	DESCRIPTION
<code>content</code>	A list of 'Form row', 'Form group' or nodes. Nodes will result in model driven display, 'Form row' and 'Form group' allow to display custom content.

See also:

- [Form row](#)
- [Form group](#)

Form definition

Actual definition of a form.




Parameters

INPUT	DESCRIPTION
<code>toolbar</code>	Definition of the toolbar displayed on top of the form.
<code>action bar</code>	Definition of the action bar displayed below the form.
<code>help URL</code>	URL of the help page.

INPUT	DESCRIPTION
<code>content</code>	The content of the layout.

Form group

Inside a 'Form', displays a collapsible group of items.



Parameters


INPUT	DESCRIPTION
<code>node</code>	If set, the group label, description and content will be model driven.
<code>label</code>	The label of the group. May override the model driven label if a node and a label are set.
<code>description</code>	The description of the group. May override the model driven description if a node and a description are set.
<code>open?</code>	Indicates if the group is initially open.
<code>content</code>	A list of 'Form row', 'Form group' or nodes. Nodes will result in model driven display, 'Form row' and 'Form group' allow to display custom content.

See also:

- [Form](#)
- [Form row](#)

Form row

Inside a 'Form' or a 'Form group', displays a row, that is basically a label and a value.



Parameters

INPUT	DESCRIPTION
<code>node</code>	If set, the label, description and content will be model driven.
<code>display_label?</code>	Indicates if the label should be displayed. If set to false in a 'Form' the content will span over both the label and content area.
<code>label</code>	The label of the row. May override the model driven label if a node and a label are set.
<code>description</code>	The description of the row. May override the model driven label if a node and a label are set.
<code>content</code>	The content to be displayed next to the label.

See also:

- [Form](#)
- [Form group](#)

Form parameter

The 'record' is the node representing the displayed record. 'current page' is a number (1 by default) representing the displayed page. This value can be changed by the 'Previous' and 'Next' buttons.

From...get

Extract a value from an object or a node.

<from "node" get "path/to/value"> is equivalent to <var "node/path/to/value">

Parameters

INPUT	DESCRIPTION
from	Base object or node, from which to extract data.
get	Path of the data to extract, inside the object in 'from'.

Function

The definition of an anonymous function. A function makes use of its parameters in its body to return a result.

Parameters

INPUT	DESCRIPTION
parameters	List of 'param' blocks, defining the parameters accepted by this function.

INPUT	DESCRIPTION
body	Body describing the result of the function, using its parameters.

See also:

- [Function parameter](#)
- [Function call](#)

Function call

Calls a function by setting actual values to its parameters.



Parameters

INPUT	DESCRIPTION
function	The function to call.
args	The arguments to pass to the function.

See also:

- [Function call argument](#)
- [Function](#)

Function call argument

Actual value to pass to the parameter of the called function.

paramName : paramValue

Parameters

INPUT	DESCRIPTION
name	The name of the parameter.
value	The actual value to pass the the parameter.

See also:

- Function call

Function parameter

This is the declaration of a parameter of a function.

param optionalParam is optional has default value 0

Parameters


INPUT	DESCRIPTION
name	The name of the parameter.
mandatory?	Indicates if this parameter must be provided by the caller.
default value	If set, this is the default value taken by this parameter if the caller does not provide one.

See also:

- Function

Grid

A layout allowing to place components in a grid with coordinates.



Parameters

INPUT	DESCRIPTION
item spacing	Space between elements, in pixels.
content	A list of 'Grid element's, which specify where to display their content.

See also:

- [Grid element](#)

Grid element

Inside a 'Grid', specifies the location of its content.

Parameters

INPUT	DESCRIPTION
row	The row where the element starts. Minimum value is 1.
column	The column where the element starts. Minimum value is 1.


INPUT	DESCRIPTION
<code>width</code>	The width of the element, in grid cells. Minimum value is 1.
<code>height</code>	The height of the element, in grid cells. Minimum value is 1.
<code>vertical alignment</code>	The vertical alignment of the element. If set, accepted values are: 'start', 'end', 'center'.
<code>horizontal alignment</code>	The horizontal alignment of the element. If set, accepted values are: 'start', 'end', 'center'.
<code>content</code>	The content inside the cell.

See also:

- [Grid](#)

If

A block returning either the content of 'then', or the content 'else', based on the result of the test.



Parameters

INPUT	DESCRIPTION
<code>test</code>	A test determining which value to return. This must be a boolean value.
<code>then</code>	The value returned if the test is true.
<code>else</code>	The value returned if the test is false.

Input parameter

Returns the raw value of an input parameter.

input parameter paramName

Parameters

INPUT	DESCRIPTION
name	The parameter name.

In role?

Returns a boolean indicating if the current user has the requested role.

in role? [administrator]

Label

Displays the label of the given node.

Label [node] var record/firstName

Parameters

INPUT	DESCRIPTION
node	The node for which to display the label.

Localized

Block which result depends on the user locale.


Localized [English (United States)] Hello [French (France)] Bonjour

Parameters

INPUT	DESCRIPTION
[locale]	Each parameter relates to a locale. The result of the block will be the one corresponding to the locale of the current user.

Message

Formats the given text by replacing java-like {0}, {1}, etc. placeholders by the argument at the given index.



Parameters

INPUT	DESCRIPTION
pattern	Text with placeholders ({0}, {1}, etc.) which will be replaced by the given arguments.
arguments	The list of values which will replace the placeholders in the pattern.

Next button

A button which will increment the 'current page' value by one.




Parameters

INPUT	DESCRIPTION
label	If set, overrides the default label.

INPUT	DESCRIPTION
default?	Indicates if this is the action to trigger when pressing 'Enter'. Only one button should be the default one.

Node

Overrides the definition of an existing node.




Parameters

INPUT	DESCRIPTION
base	The node to override.
mandatory?	If set, overrides the mandatory indicator of the base node.
label	If set, overrides the label of the node.
description	If set, overrides the label of the node.
widget	If set, overrides the default widget of the node.

Nomenclature item

A nomenclature item is a (key, label) pair.




Parameters

INPUT	DESCRIPTION
value	The value of the item.
label	The label of the item.

Not

Returns the inverse of the given boolean.



Object

An object composed of a list of (key, value) pairs.



Parameters


INPUT	DESCRIPTION
content	A list of 'property' blocks.

See also:

- [Property](#)

Otherwise

Final choice in a 'choose' block. Will be returned if no 'when' case matched.



Parameters


INPUT	DESCRIPTION
<code>content</code>	The value to return if no 'when' case matched.

See also:

- [Choose](#)

Paragraph title

Display a paragraph title.




Parameters

INPUT	DESCRIPTION
<code>label</code>	The text to display.
<code>horizontal alignment</code>	The horizontal alignment of the element. If set, accepted values are: 'start', 'end', 'center'.

Password

Displays the input password widget.




Parameters

INPUT	DESCRIPTION
node	The node for which to display the widget.
readonly?	Indicates if the widget should be in readonly mode.

Previous button

A button which will decrement the 'current page' value by one.



Parameters

INPUT	DESCRIPTION
label	If set, overrides the default label.
default?	Indicates if this is the action to trigger when pressing 'Enter'. Only one button should be the default one.

Property

A property of an object.


property : value

Parameters

INPUT	DESCRIPTION
name	Name of the property. Must be unique for a given object.
value	Value of the property.

Radio button

Displays a radio button widget.




Parameters

INPUT	DESCRIPTION
node	The node for which to display the widget.

INPUT	DESCRIPTION
<code>readonly?</code>	Indicates if the widget should be in readonly mode.
<code>label enabled?</code>	If set, indicates if the item label is to be added next to the widget.
<code>index</code>	The index for this enumeration item.
<code>specific nomenclature</code>	If set, overrides the model-driven nomenclature.

Radio button group

Displays the radio button group widget.



Parameters

INPUT	DESCRIPTION
<code>node</code>	The node for which to display the widget.

INPUT	DESCRIPTION
<code>readonly?</code>	Indicates if the widget should be in readonly mode.
<code>columns number</code>	If set, defines the number of columns to use for the layout of the radio buttons.
<code>specific nomenclature</code>	If set, overrides the model-driven nomenclature.

Revert button

Standard 'Revert' button.



Parameters

INPUT	DESCRIPTION
<code>label</code>	If set, overrides the default label.
<code>default?</code>	Indicates if this is the action to trigger when pressing 'Enter'. Only one button should be the default one.

Save button

Standard 'Save' button.




Parameters

INPUT	DESCRIPTION
<code>label</code>	If set, overrides the default label.

INPUT	DESCRIPTION
default?	Indicates if this is the action to trigger when pressing 'Enter'. Only one button should be the default one.

Save and close button

Standard 'Save and close' button.




Parameters

INPUT	DESCRIPTION
label	If set, overrides the default label.
default?	Indicates if this is the action to trigger when pressing 'Enter'. Only one button should be the default one.

Tab

A tab, inside a 'Tab container'.



Parameters

INPUT	DESCRIPTION
node	If present, the tab title, description and content will be model driven.


INPUT	DESCRIPTION
<code>title</code>	The title of the tab. May override the model driven title if a node and a title are set.
<code>description</code>	The description of the tab. May override the model driven description if a node and a description are set.
<code>content</code>	The content of the tab. May override the model driven content if a node and a content are set.

See also:

- [Tab container](#)

Tab container

A container of tabs, which displays a list of tab names and one active tab.



Parameters


INPUT	DESCRIPTION
<code>content</code>	A list of 'Tab's.

See also:

- [Tab](#)

Text input

Displays the text input widget.




Parameters

INPUT	DESCRIPTION
node	The node for which to display the widget.
readonly?	Indicates if the widget should be in readonly mode.
height	If set, defines the height of the widget, in pixels.
width	If set, defines the width of the widget, in pixels.

INPUT	DESCRIPTION
<code>multiline?</code>	If set, defines if the widget spans over multiple lines.
<code>background</code>	If set, defines the background color, in hexadecimal format.
<code>foreground</code>	If set, defines the foreground color, in hexadecimal format.

Toolbar named

The toolbar having the specified name.




Parameters

INPUT	DESCRIPTION
<code>name</code>	The name of the toolbar to display.

Variable

A reference to a variable declared by a 'with...declare' block. A path is also accepted if the variable is an object or a node.




Parameters

INPUT	DESCRIPTION
<code>expression</code>	A reference to a variable declared by a 'with...declare' block. A path is also accepted if the variable is an object or a node.

When

A choice in a 'choose' block. If the test is true, the 'then' content will be returned, otherwise the next block will be tested.



Parameters


INPUT	DESCRIPTION
when	A boolean test indicating if this case should be resolved.
then	If the test is true, this value will be returned.

See also:

- [Choose](#)

Widget

Displays the default widget.




Parameters

INPUT	DESCRIPTION
node	The node for which to display the widget.
readonly?	Indicates if the widget should be in readonly mode.

With

Declares variables that can be used in its 'return' statement.



Parameters

INPUT	DESCRIPTION
variables	A list of 'declare' blocks.
return	The value to return. The variables declared in the 'variables' statement can be used here, with 'var' blocks.

See also:

- [Declare](#)

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.

- Bean categories
- Sample of ScriptTask
 - Java Code
- Sample of ScriptTaskBean
 - Java Code
 - Configuration through module.xml
- Samples of UserTask
 - Service declaration via module.xml
 - A more complex UserTask
- Samples of Condition
 - Java Code
- Sample of ConditionBean
 - Java Code
 - Configuration through module.xml
- Sample of SubWorkflowsInvocationBean
 - Java Code
 - Configuration through module.xml
- Sample of WaitTaskBean
 - Java Code
 - Configuration through module.xml
- Sample of ActionPermissionsOnWorkflow
 - Java Code
 - Configuration through module.xml
- Sample of WorkflowTriggerBean
 - Java Code

- Configuration through module.xml
- Sample of trigger starting a process instance
 - Sample

Bean categories

STEP	LIBRARY	SPECIFIC
Scripts	ScriptTaskBean	ScriptTask
Conditions	ConditionBean	Condition
User task	UserTask	

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
    RuntimeException
    {
        Repository repository = aContext.getRepository();

        String initialBranchString =
        aContext.getVariableString("initialBranch");

        AdaptationHome initialBranch =
        repository.lookupHome(HomeKey.forBranchName(initialBranchString));

        if (initialBranch == null)

            throw RuntimeException.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](#)

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`

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>

```

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](#)

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](#)

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](#)

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>

    </bean>

  </beans>

</module>
```



```

        </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>

```

Sample of SubWorkflowsInvocationBean

Java Code

```
public class MySubWorkflowsInvocationBean extends
SubWorkflowsInvocationBean

{

    @Override

    public void handleCreateSubWorkflows(SubWorkflowsCreationContext
aContext)

        throws RuntimeException

    {

        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 RuntimeException

    {

        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").e
quals("true")

            &&

```

```

validateProductDirection.getDataContext().getVariableString("Accepted").equals(

        "true"))

        aContext.setVariableString("validation", "ok");

    }

}

```

See also:

- [com.orchestranetworks.workflow.SubWorkflowsInvocationBean](#)

Configuration through module.xml

`SubWorkflowsInvocationBean` bean must be declared in `module.xml`:

```

<module>

    <beans>

        <bean
className="com.orchestranetworks.workflow.test.MySubWorkflowsInvocationBean"/>

    </beans>

</module>

```

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](#)

Configuration through module.xml

`WaitTaskBean` bean must be declared in `module.xml` :

```
<module>

    <beans>

        <bean
className="com.orchestranetworks.workflow.test.MyWaitTaskBean"/>

    </beans>

</module>
```

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](#)

Configuration through module.xml

`ActionPermissionsOnWorkflow` bean must be declared in `module.xml` :


```

<module>

    <beans>

        <bean
className="com.orchestranetworks.workflow.test.MyDynamicPermissions"/>

    </beans>

</module>

```

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.T0, "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 RuntimeException

    {

        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.T0, "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.T0, "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.T0, "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](#)

Configuration through module.xml

`WorkflowTriggerBean` bean must be declared in `module.xml` :

```
<module>

    <beans>

        <bean
className="com.orchestranetworks.workflow.test.MyWorkflowTriggerBean"/>

    </beans>

</module>
```

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

        WorkflowEngine engine =
        WorkflowEngine.getFromProcedureContext(aContext.getProcedureContext());

        ProcessLauncher launcher =
        engine.getProcessLauncher(processPublishedKey);

        //initialize Data Context

        launcher.setInputParameter("code", "/root/Client[./code=\"\" + code
+ "\"]");

        launcher.setInputParameter("workingBranch",
aContext.getAdaptationHome().getKey().getName());

        //Starts process

        launcher.launchProcess();

    }

    //...

```

}

User interface

- [Interface customization](#)
- [User services](#)
 - [Overview](#)
 - [Quick start](#)
 - [Implementing a user service](#)
 - [Declaring a user service](#)
- [Development recommendations](#)

Interface customization

The TIBCO EBX® graphical interface can be customized through various EBX® APIs.

- [Using EBX® as a Web Component](#)
- [Adding user services](#)
- [Customizing forms](#)
- [Customizing widgets](#)
- [Customizing table filter](#)
- [Customizing record label](#)
- [Including CSS and JavaScript](#)

Using EBX® as a Web Component

EBX® can be integrated into any application that is accessible through a [supported web browser](#), thanks to the Web Component API.

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](#).

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](#)

Adding 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](#)

Customizing forms

It is possible to override the default layout and behavior of forms in the user interface by using various tools and API.

Custom forms editor

It is possible to override the default layout of forms in the user interface by using the 'Custom forms' extension in the DMA. This extension provides a graphical editor, which gives 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:

- [Custom forms](#)

Programmatic form layout	It is also possible to use the <code>UIForm</code> API to override the default layout of forms. 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: • <code>UIForm</code> • <code>UIFormPaneWriter</code> • <code>UIWidget</code> • <code>UIFormHeader</code> • <code>UIFormBottomBar</code>
User service as form layout	It is also possible to use user services to override the default layout of forms. This API gives the possibility to customize the layout of a form, while having the same components as record forms using the default layout, with a customizable behavior. See also: • <code>UserServiceRecordFormFactory</code> • Overview

Customizing 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`

Customizing 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`
- [Properties of data model elements](#)
- `UILabelRendererForHierarchy`

Customizing 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`

Including 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](#)
- [Development recommendations](#)
- `UIDependencyRegisterer`

User services

- [Overview](#)
- [Quick start](#)
- [Implementing a user service](#)
- [Declaring a user service](#)

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.

- [Nature](#)
- [Declaration](#)
- [Display](#)
- [Legacy user services](#)

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

NATURE	DESCRIPTION
	the current selection implies selecting a <i>single</i> association target record.

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.

Display

On the following figure are displayed the functional areas of a user service.

The screenshot shows a user service form for 'John Doe'. The form is divided into several sections, with numbered callouts indicating functional areas:

- 1**: User identification area (User : John Doe).
- 2**: Help icon (?) and Close icon (X).
- 3**: Actions dropdown menu.
- 4**: Navigation buttons (Previous < and Next >).
- 5**: Tabbed interface (Main and Groups).
- 6**: Vertical scrollbar.
- 7**: Form fields for personal information (Last Name, First Name, Start Date, Country, Job, Email, Phone, Mobile Phone).
- 8**: Action buttons (Save, Save and close, Revert, Close).
- 9**: Footer area (SkypeId, jdoe).

A. Header	1. Breadcrumb
B. Form	2. Message box button
	3. Top toolbar
	4. Navigation buttons
	5. Help button
	6. Close button (pop-ups only)
	7. Tabs
	8. Form pane (one per tab)
	9. Bottom buttons

Most areas are optional and customizable. Refer to [Quick start](#), [Implementing a user service](#) and [Declaring a user service](#) for more details.

Legacy user services

Before the 5.8.0 version, user services were declared in XML and based on Servlet/JSP. Although this type of declaration should no longer be used, the [legacy documentation](#) is still available.

Quick start


- [Main classes](#)
- [Hello world](#)

Main classes

The minimum requirement is to implement two classes, one for the service declaration and one for the implementation itself.

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 instanciated 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 [SchemaExtensions](#).

Implementing a user service

- [Implementation interface](#)
- [Life cycle and threading model](#)
- [Object Context](#)
- [Display setup](#)
- [Database updates](#)
- [Ajax](#)
- [REST data services](#)
- [File upload](#)
- [File download](#)
- [User service without display](#)
 - [Known limitation](#)

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>

NATURE	INTERFACE
Association	UserService<AssociationEntitySelection>
AssociationRecord	UserService<AssociationRecordEntitySelection>

Life cycle and threading model

The user service implementation class is:

- Instantiated at the first HTTP request by a call to its declaration `createUserService()` 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.

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.setupObjectContext .
```

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](#), [path](#) 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 static final ObjectKey _PersonObjectKey =
ObjectKey.forName("person");

    private static final ObjectKey _PartnerObjectKey =
ObjectKey.forName("partner");
```

```

// Define a path for each property:

private static final Path _FirstName = Path.parse("firstName");

private static final Path _LastName = Path.parse("lastName");

private static final 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);

    }

    ...

}

```

Display setup

The display is set up by the user service's implementation of the method

`UserService.setupDisplay`.

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.setupDisplay`.

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 `UserServicePane` to render a plain EBX® form. The callback can also be an instance of `UserServiceTabbedPane` to render an EBX® form with tabs.

For specific cases, the callback can implement `UserServiceRawPane`. 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 static final 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 ULR 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`, 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` 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;

}

}
```

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 static final ObjectKey _Person1_ObjectKey =
ObjectKey.forName("person1");

    private static final 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](#):

```

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;

}

```

Ajax

A user service can implement Ajax callbacks. An Ajax callback must implement the interface `UserServiceAjaxRequest`.

The client calls an Ajax callback using the URL generated by:

```
UserServiceResourceLocator.getURLForAjaxRequest .
```

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`.

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>");

    }

}

```

REST data services

A user service can access REST data services through HTTP requests.

The client should use the URL generated by: `UIResourceLocator.getURLForRest`. This URL includes required information for the user authentication.

For more information on REST data services see the [Built-in RESTful services](#).

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-dataservices/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("{}");

    }

}

```

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 static final ObjectKey _File_ObjectKey =
ObjectKey.forName("file");


    // Paths for the "file" object.

    public static final Path _Title = Path.parse("title");

    public static final 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 `UIWidgetFileUpload`.

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.getDateTimeInstance(

            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();

    }

}

```

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](#) and method

`UserServiceExtended.initialize`.

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.

Declaring a user service

- [Declaration interface](#)
- [Life cycle and threading model](#)
- [Registration](#)
- [Service properties](#)
- [Service activation scope](#)
- [Web component declaration](#)

- Parameters declaration and availability in workflows and perspectives
- User service extension
- User service groups
 - Built-in User Service Groups
 - Declaring a User Service Group
 - Associating a service to a group

Declaration interface

The following table lists, per nature, the interface that the declaration class of a user service must implement:

NATURE	DECLARATION INTERFACE
Dataspace	<code>UserServiceDeclaration.OnDataspace</code>
Dataset	<code>UserServiceDeclaration.OnDataset</code>
TableView	<code>UserServiceDeclaration.OnTableView</code>
Record	<code>UserServiceDeclaration.OnRecord</code>
Hierarchy	<code>UserServiceDeclaration.OnHierarchy</code>
HierarchyNode	<code>UserServiceDeclaration.OnHierarchyNode</code>
Association	<code>UserServiceDeclaration.OnAssociation</code>
AssociationRecord	<code>UserServiceDeclaration.OnAssociationRecord</code>

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.

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](#) and `ModuleRegistrationServlet`.

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`.

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`.

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.includeAllDataspace(DataspaceType.BRANCH);

        aContext.excludeDataspaceMatching(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.defineActivation`.

For more information about the resolution of the user service availability, see [Resolving permissions on services](#).

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.declareWebComponent` 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](#) 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](#) and `ModuleRegistrationServlet`.

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`.

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](#) provided by EBX®,
- [Custom User Service Groups](#) declared in a module.

The link between groups and services is made upon service declaration. See [Associating a service to a group](#).

Built-in User Service Groups

Available built-in service groups:

Service Group Key	Description
<code>@ebx-importExport</code>	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.
<code>@ebx-views</code>	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` :

```
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](#), using the method

`UserServicePropertiesDefinitionContext.setGroup` :

```

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.

Development recommendations

- [HTML](#)
 - [XHTML](#)
 - [iFrames](#)
- [CSS](#)
 - [Public CSS classes](#)
 - [Advanced CSS](#)
- [JavaScript](#)
 - [Public JS functions](#)
 - [JavaScript call during page generation in Java](#)
 - [Advanced JavaScript](#)

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

`UIHttpManagerComponent`. 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 `UIComponentWriter`, uses a

`UIHttpManagerComponent` 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("\" src=\"\" >").add("</
iframe>");

// launch the iFrame from JavaScript

writer.addJS("document.getElementById(\"").addJS(iFrameId).addJS("\").src
= \"").addJS(iFrameURL).addJS("\";");
```

CSS

Public CSS classes

The constant catalog `UICSSClasses` 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`

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](#).

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.

ebx_	Internal built-in
yui	Yahoo User Interface global
ygtv	Yahoo User Interface tree view
layout-doc	Yahoo User Interface layout
cke_	CK editor (used by HTML editor widget)
fa	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®.

selected	YUI selected tree node
hd	YUI floating pane header
bd	YUI floating pane body
ft	YUI floating pane footer
container-close	YUI inner popup close button
underlay	YUI inner popup shadow
hastitle	YUI menu group with title

topscrollbar	YUI menu top scroll zone
bottomscrollbar	YUI menu bottom scroll zone
withtitle	YUI calendar
link-close	YUI calendar close button
collapse	YUI layout closed pane indicator
pull-right	Font Awesome parameter
pull-left	Font Awesome parameter

Examples to avoid conflicts

Don't

```
.selected {
    background-color: red;
}
```

Do

```
#myCustomComponent li.selected {
    background-color: red;
}
```

JavaScript

Public JS functions

The catalog of JavaScript functions [JavaScriptCatalog](#) 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 `UIJavaScriptWriter`.

This JavaScript is executed once the whole page is loaded. It is possible to instantly manage the HTML elements written with `UIBodyWriter.add`. 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](#). 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 `UIDependencyRegisterer` 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.

ebx_	Internal built-in API
EBX_	Internal built-in API
YAHOO	Yahoo User Interface API

SOAP data services

- [Introduction](#)

- WSDL generation
- SOAP operations

Introduction

- Overview
- Activation and configuration
- Interactions
 - Input and output message encoding
 - Tracking information
 - Session parameters
 - Exception handling
 - Using JMS
- Security
 - Authentication
 - Global permissions
 - Overriding the SOAP security header
 - Lookup mechanism
- Monitoring
- SOAP and REST comparative
- Limitations
 - Date, time & dateTime format
 - SOAP naming convention

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](#), for an integration use case, a [WSDL](#) 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](#).

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](#) for more information.

In case of specific deployment, for example using reverse-proxy mode, see [URLs computing](#) 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](#) and [Using JMS](#) for more information.

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.getTrackingInfo` 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.getInputParameterValue` 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](#). 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`.

Security

Authentication

Authentication is mandatory to access to the content. 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 QWxpYmFiYTpvcmGVuIHNIc2FtZQ==
```

- '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](#) section.

For more information on this authentication scheme, see

`Directory.authenticateUserFromLoginPassword`.

- 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/
secect">

    <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](#).

- '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.authenticateUserFromHttpRequest` for more information.
- The 'SOAP Specific Header Authentication Scheme'.

For more information, see [Overriding the SOAP security header](#).

Global permissions

Global access permissions can be independently defined for the SOAP and WSDL connector accesses.

See also:

- [Global permissions](#)

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:

Schema location	The URI of the Security XML Schema to import into the WSDL.
Target namespace	The target namespace of elements in the schema.
Namespace prefix	The prefix for the target namespace.
Message name	The message name to use in the WSDL.

Root element name	The root element name of the security header. The name must be the same as the one declared in the schema.
wSDL:part element name	The name of the <code>wSDL:part</code> 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.authenticateUserFromSOAPHeader` .

Note:

Only available for [SOAP operations](#).

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	SOAP Security Header • The SOAP request is received over the JMS protocol. • The SOAP header content must contains a <code>Security</code> element. SOAP Specific Header • The SOAP request is received over the JMS protocol. • The SOAP header content must not contain a <code>Security</code> element.
SOAP / HTTP	Basic • The HTTP request must hold an <code>Authorization</code> header. • <code>Authorization</code> header value must start with the word <code>Basic</code>. • No <code>login</code> is provided in the URL parameters. Standard • The HTTP request must not hold an <code>Authorization</code> header. • A <code>login</code> and a <code>password</code> are provided in the URL parameters. SOAP Security Header • The SOAP header content must contain a <code>Security</code> element.

OPERATION / PROTOCOL	AUTHENTICATION METHODS AND APPLICATION CONDITIONS
	- The HTTP request must not hold an <code>Authorization</code> header. - No <code>login</code> is provided in the URL parameters. Specific - The HTTP request must not satisfy the conditions of the previous authentication methods. SOAP Specific Header - The SOAP header content must not contain a <code>Security</code> element. - The HTTP request must not hold an <code>Authorization</code> header. - No <code>login</code> is provided in the URL parameters.
WSDL / HTTP	Basic - The HTTP request must not hold an <code>Authorization</code> header. <code>Authorization</code> header value must start with the word <code>Basic</code>. - No <code>login</code> is provided in the URL parameters. Standard - The HTTP request must not hold an <code>Authorization</code> header. - A <code>login</code> and a <code>password</code> are provided in the URL parameters.

OPERATION / PROTOCOL	AUTHENTICATION METHODS AND APPLICATION CONDITIONS
	Specific The HTTP request must not satisfy the conditions of the previous authentication methods.

In case of multiple authentication methods present in the same request, EBX® will return an HTTP code `401 Unauthorized`.

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](#)
- [TIBCO EBX® main configuration file](#)

SOAP and REST comparative

OPERATIONS	SOAP	REST
Data		
Select metamodel		X
Select or count records (with filter and/or view publication)	X	X
Select records with association values (first level)	X	

OPERATIONS	SOAP	REST
Update or delete records (with filter and/or view publication)		X
Selector for possible enumeration values (with filter)		X
Prepare for creation or duplication		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	
Generate service contract	WSDL	OpenAPI
Views		
Look up published table views		X
Dataspaces		

OPERATIONS	SOAP	REST
Select dataspace or snapshot information		X
Select root or children dataspace, or select snapshots		X
Create, close, merge a dataspace	X	X
Create, close a snapshot	X	X
Validate a dataspace or a snapshot	X	
Validate a dataset	X	
Locking, unlocking a dataspace	X	X
Workflow		
Start, resume or end a workflow	X	
Administration		
Manage the default directory content 'Users', 'Roles'... tables.	X	X
Open, close the user interface	X	X
Select, insert, update, delete operations for administration dataset		X
Select the system information	X	X
Staging API		X (*)

OPERATIONS	SOAP	REST
Other		
Develop web services from the Java API		X (*)

(*) See [REST Toolkit](#) for more information.

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.

WSDL generation

- [Supported standard](#)
- [Operation types](#)
- [WSDL download methods](#)
 - [WSDL download from the user interface](#)
 - [WSDL download from HTTP protocol](#)
- [HTTP examples](#)

Supported standard

TIBCO EBX® generates a WSDL that complies with the [W3C Web Services Description Language 1.1](#) standard.

Operation types

A WSDL can be generated for different types of operations:

OPERATION TYPE	WSDL DESCRIPTION
<code>custom</code>	WSDL for EBX® add-ons.
<code>dataset</code>	WSDL for dataset and replication operations.
<code>directory</code>	WSDL for default EBX® directory operations. It is also possible to filter data using the <code>tablePaths</code> or <code>operations</code> parameters.
<code>repository</code>	WSDL for dataspace or snapshot management operations.
<code>tables</code>	WSDL for operations on the tables of a specific dataset.
<code>userInterface</code>	Deprecated since version 5.8.1. This operation type has been replaced by <code>administration</code>. While the user interface management operations are still available for backward compatibility reasons, it is recommended to no longer use this type. WSDL for user interface management operations (these operations can only be accessed by administrators).
<code>administration</code>	WSDL for administration operations like: • user interface management • system information retrieval

OPERATION TYPE	WSDL DESCRIPTION
	These operations can only be accessed by administrators.
<code>workflow</code>	WSDL for EBX® workflow management operations.

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
<code>custom</code>	All profiles, if at least one web service is registered.
<code>dataset</code>	All profiles.
<code>directory</code>	All profiles, if the following conditions are valid: • No specific directory implementation is used. (The built-in Administrator role is only subject to this condition). • Global access permissions are defined for the administration. • 'Directory' dataset permissions have writing access for the current profile.
<code>repository</code>	All profiles.
<code>tables</code>	All profiles.

OPERATION TYPE	ACCESS RIGHT PERMISSIONS
userInterface	Deprecated since version 5.8.1. This operation type has been replaced by <code>administration</code>. While the user interface management operations are still available for backward compatibility reasons, it is recommended to no longer use this type. Built-in administrator role or delegated administrator profiles, if all conditions are valid: • Global access permissions are defined for the administration. • 'User interface' dataset permissions have writing access for the current profile.
administration	Built-in administrator role or delegated administrator profiles, if all conditions are valid: • Global access permissions are defined for the administration. • 'Administration' dataset permissions have write access for the current profile.
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](#) 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](#).
- **<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. Empty value.

PARAMETER NAME	PATHPARAM	QUERYPARAM	REQUIRED	DESCRIPTION
login	no	yes	no	Specifies the user identifier. Required when the standard authentication method is used. String type value.
password	no	yes	no	Specifies the user password. Required when the standard authentication method is used. String type value.
type	yes	no	yes	Specifies the operation type. Possible values are: custom , dataset , directory , administration , userInterface , repository ,

PARAMETER NAME	PATHPARAM	QUERYPARAM	REQUIRED	DESCRIPTION
				<code>tables</code> or <code>workflow</code>. <code>String</code> type value.
<code>branch</code> <code>version</code>	yes	yes	(*)	Specifies the dataspace or the snapshot identifier. (*) required for <code>tables</code> and <code>dataset</code> types, ignored otherwise. <code>String</code> type value.
<code>instance</code>	yes	yes	(*)	Specifies the dataset name. <code>String</code> type value.
<code>tablePaths</code>	no	yes	no	Specifies the list of table paths. Optional for <code>tables</code> or <code>directory</code> types, ignored otherwise.

PARAMETER NAME	PATHPARAM	QUERYPARAM	REQUIRED	DESCRIPTION
				If not defined, all tables are selected. Each table path is separated by a comma character. <code>String</code> type value.
<code>operations</code>	no	yes	no	Allows generating a WSDL for an operations subset. Optional for <code>tables</code> or <code>directory</code> operation types, ignored otherwise. If not defined, all operations for the given type are generated. This parameter's value is a concatenation of one or more

PARAMETER NAME	PATHPARAM	QUERYPARAM	REQUIRED	DESCRIPTION
				of the following characters: • C = Count record(s) • D = Delete record(s) • E = Get credentials • G = Get changes • I = Insert record(s) • U = Update record(s) • R = Read operations (equivalent to CEGS) • S = Select record(s) • W = Write operations (equivalent to DIU) <code>String</code> type value.
<code>namespaceURI</code>	yes	yes	(**)	Specifies the unique name space URI of

PARAMETER NAME	PATHPARAM	QUERYPARAM	REQUIRED	DESCRIPTION
				the custom web service. (**)Is required when <code>type</code> parameter is set to <code>custom</code> , ignored otherwise. <code>URI</code> type value.
<code>attachmentFilename</code>		yes	(***)	Specifies the attachment file name. (***) optional if <code>isContentInAttachment</code> parameter is defined and set to <code>true</code> , ignored otherwise. <code>String</code> type value.
<code>isContentInAttachment</code>		yes	no	Specifies if the WSDL is downloaded as an attachment. <code>Boolean</code> type value.

PARAMETER NAME	PATHPARAM	QUERYPARAM	REQUIRED	DESCRIPTION
				Default value is <code>false</code> .
<code>targetNamespace</code> no		yes	no	Overrides the target namespace URI of the WSDL. <code>URI</code> type value. Default value is <code>urn:ebx:ebx-dataservices</code>.

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
<code>200</code> (OK)	The WSDL content was successfully generated and is returned by the request (optionally in an attachment).
<code>400</code> (Bad request)	The request is incorrect. This occurs when: • A request element is incorrect. • The unicity check on table names contains at least one error.

STATUS CODE	INFORMATION
	Note: See WSDL and table operations for more information.
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`.

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`

SOAP operations

- Operations generated from a data model
 - Content policy
 - Common request parameters
 - Select operations
 - Delete operation
 - Count operation
 - Update operation
 - Insert operation
 - Get changes operations
 - Get credentials operation
 - Multiple chained operations
 - Optimistic locking
- Operations on datasets and dataspaces
 - Validate a dataspace
 - Validate a dataset
 - Create a dataspace
 - Create a snapshot
 - Locking a dataspace
 - Unlocking a dataspace
 - Merge a dataspace
 - Close a dataspace or snapshot
 - Replication refresh
- Operations on data workflows
 - Start a workflow

- Resume a workflow
- End a workflow
- Administrative services
 - Directory services
 - User interface operations
 - System information operation

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](#) 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](#)
- [Permissions](#)

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.

ELEMENT	DESCRIPTION	REQUIRED
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	XPath predicate 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. See also: <code>includesMetadata</code> parameter	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 Import .	Only required for the insert and update operations
viewPublication	This parameter can be combined with the predicate parameter as a logical <code>AND</code> operation.	No

ELEMENT	DESCRIPTION	REQUIRED
	The behavior of this parameter is described in the section EBX® as a Web Component. It cannot be used if the 'viewId' parameter is used, and cannot be used on hierarchical nor tile views.	
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. This parameter cannot be used if the 'viewPublication' parameter is used.	No
blockingConstraintsDisabled	This property is available for all table updates data service operations. If <code>true</code>, the validation process disables blocking constraints defined in the data model. If this parameter is not present, the default is <code>false</code>.	No

ELEMENT	DESCRIPTION	REQUIRED
	See Blocking and non-blocking constraints for more information.	
details	The <code>details</code> element specifies the following option: The optional attribute <code>locale</code> (default 'en-US') defines the language of the blockingConstraintsDisabled parameter in which the validation messages must be returned.	No
disableRedirectionToLastBroadcast	This property is available for all data service operations. 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 snapshot broadcast. If this parameter is not present, the default is <code>false</code> (redirection on a D3 master enabled), unless the configuration property ebx.dataservices.disableRedirectionToLastBroadcast.default has been set.	No

ELEMENT	DESCRIPTION	REQUIRED
	If the specified dataspace is not a delivery dataspace on a D3 primary node, this parameter is ignored.	

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>

    <includesMetadata>String</includesMetadata>

    <viewPublication>String</viewPublication>

    <exportCredentials>boolean</exportCredentials>

    <pagination>

        <previousPageLastRecordPredicate>String</
previousPageLastRecordPredicate>

        <pageSize>Integer</pageSize>

    </pagination>

    <disableRedirectionToLastBroadcast>boolean</
disableRedirectionToLastBroadcast>

</m:select_{TableName}>

```


with:

ELEMENT	DESCRIPTION	REQUIRED
branch	See the description under Common parameters .	
version	See the description under Common parameters .	
instance	See the description under Common parameters .	
predicate	See the description under Common parameters. This parameter can be combined with the viewPublication parameter as a logical AND operation.	
viewPublication	See the description under Common parameters .	
includesMetadata	The response contains the specified metadata. Each returned record contains additional elements called steps under the ebx-metadata root element. For instance, a 'system' step value will contain technical data. See also the optimistic locking section: <pre><system> <uuid>7FE03810-6A67-11ED-A892-00FF20216100</uuid> <creator>Uadmin</creator></pre>	No

ELEMENT	DESCRIPTION	REQUIRED
	<pre><creation_time>2022-11-12T00:07:40.858</creation_time></pre> <pre><updater>Uadmin</updater></pre> <pre><update_time>2022-11-12T00:07:40.831</update_time></pre> <pre></system></pre> String value steps (default value is empty) must be separated by a comma: <code>system, teamUp</code>. All steps can be returned using the <code>ebx-all</code> value.	
includesTechnicalData	<i>Deprecated since version 6.1.0 and replaced by <code>includesMetadata</code> request parameter.</i> The response will contain technical data if <code>true</code>. See also the optimistic locking section. Each returned record will contain additional attributes for this technical information, for instance: <pre>...<table</pre> <pre>ebxdlastTime="2010-06-28T10:10:31.046" ebxdlastUser="Uadmin</pre> <pre>ebxduuid="9E7D0530-828C-11DF-B733-0012D01B6E76">... .</pre>	No

ELEMENT	DESCRIPTION	REQUIRED
exportCredentials	If <code>true</code> the select will also return the credentials for each record.	No
pagination	Specifies the pagination configuration, see child elements below.	No
pageSize (nested under the <code>pagination</code> element)	Specifies the maximum number of records per page. <code>Integer</code> type, default value is 10. See also: ebx.dataservices.pagination.pageSize.default The value 0 selects the maximum recommended page size, default is 10000. See also: ebx.dataservices.pagination.maxSize.default Note: If the specified or default page size value exceeds the maximum recommended page size, the maximum page size is used.	No
previousPageLastRecordPredicate (nested under the <code>pagination</code> element)	When pagination is enabled, XPath predicate that defines the record after which the page must fetched, this value is provided by the	No

ELEMENT	DESCRIPTION	REQUIRED
	previous response, as the element <code>lastRecordPredicate</code> . If the passed record is not found, the first page will be returned.	
<code>disableRedirectionToLastBroadcast</code>	See the description under Common parameters .	

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>

            <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](#) 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 Common parameters .	
version	See the description under Common parameters .	
instance	See the description under Common parameters .	
disableRedirectionToLastBroadcast	See the description under Common parameters .	

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>

            </XX>

        </settings>

    </data>

</ns1:selectInstanceResponse>

```

```

        <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>

    <instance>String</instance>

    <predicate>String</predicate>

    <includeOcculting>boolean</includeOcculting>

    <inheritIfInOccultingMode>boolean</inheritIfInOccultingMode>

    <checkNotChangedSinceLastTime>dateTime</checkNotChangedSinceLastTime>

    <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>

    <details locale="Locale"/>

```

```

    <disableRedirectionToLastBroadcast>Boolean</
disableRedirectionToLastBroadcast>

</m:delete_{TableName}>

```

with:

ELEMENT	DESCRIPTION	REQUIRED
branch	See the description under Common parameters .	
instance	See the description under Common parameters .	
predicate	See the description under Common parameters .	
includeOcculting	Includes the records in occulting mode. Default value is <code>false</code> .	No
inheritIfInOccultingMode	Inherits the record from its parent if it is in occulting mode. Default value is <code>false</code> .	No
occultIfInherit	<i>Deprecated since version 5.7.0</i> Occults the record if it is in inherit mode. Default value is <code>false</code> .	No
checkNotChangedSinceLastTime	Timestamp used to ensure that the record has not been modified since the last read. Also see the optimistic locking section.	No
blockingConstraintsDisabled	See the description under Common parameters .	

ELEMENT	DESCRIPTION	REQUIRED
details	See the description under Common parameters .	
disableRedirectionToLastBroadcast	See the description under Common parameters .	

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. '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>

    <disableRedirectionToLastBroadcast>boolean</
disableRedirectionToLastBroadcast>

</m:count_{TableName}>
```

with:

ELEMENT	DESCRIPTION
branch	See the description under Common parameters .
version	See the description under Common parameters .
instance	See the description under Common parameters .
predicate	See the description under Common parameters .
disableRedirectionToLastBroadcast	See the description under Common parameters .

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>

    <instance>String</instance>

    <updateOrInsert>boolean</updateOrInsert>

    <byDelta>boolean</byDelta>

    <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>

    <details locale="Locale"/>

    <disableRedirectionToLastBroadcast>boolean</
disableRedirectionToLastBroadcast>

    <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 Common parameters .	
instance	See the description under Common parameters .	
updateOrInsert	If <code>true</code> and the record does not currently exist, the operation creates the record. <code>boolean</code> type, the default value is <code>false</code> .	No
byDelta	If <code>true</code> and an element does not currently exist in the incoming message, the target value is not changed. If <code>false</code> and node is declared <code>hiddenInDataServices</code> , the target value is not changed. The complete	No

ELEMENT	DESCRIPTION	REQUIRED
	behavior is described in the sections Insert and update operations .	
blockingConstraintsDisabled	See the description under Common parameters .	
details	See the description under Common parameters .	
disableRedirectionToLastBroadcast	See the description under Common parameters .	
data	See the description under Common parameters .	

See also:

- [Optimistic locking](#)

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. '95' indicates that at least one operation has violated a blocking

ELEMENT	DESCRIPTION
	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>

  <instance>String</instance>

  <byDelta>boolean</byDelta>

  <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>

  <details locale="Locale"/>

  <disableRedirectionToLastBroadcast>boolean</
disableRedirectionToLastBroadcast>

  <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 Common parameters .	
instance	See the description under Common parameters .	
byDelta	If <code>true</code> and an element does not currently exist in the incoming message, the target value is not changed. If <code>false</code> and node is declared <code>hiddenInDataServices</code> , the target value is not changed. The complete behavior is described in the sections Insert and update operations .	No
blockingConstraintsDisabled	See the description under Common parameters .	
details	See the description under Common parameters .	
disableRedirectionToLastBroadcast	See the description under Common parameters .	

ELEMENT	DESCRIPTION	REQUIRED
data	See the description under Common parameters .	

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. '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](#).

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>

    <pagination>

        <pageSize>Integer</pageSize>

    </pagination>

    <disableRedirectionToLastBroadcast>boolean</
disableRedirectionToLastBroadcast>

</m:getChanges_{TableName}>

```

with:

ELEMENT	DESCRIPTION	REQUIRED
branch	See the description under Common parameters .	
version	See the description under Common parameters .	
instance	See the description under Common parameters .	
compareWithBranch	The identifier of the dataspace with which to compare.	One of either this parameter or the ' compareWithVersion ' parameter must be defined.
compareWithVersion	The identifier of the snapshot with which to compare.	One of either this parameter or the ' compareWithBranch ' parameter must be defined.
compareWithInstance	The identifier of the dataset with which to compare. If it is undefined, instance parameter is used.	No
resolvedMode	Defines whether or not the difference is calculated in resolved mode. Default is	No

ELEMENT	DESCRIPTION	REQUIRED
	<code>true</code> . See Resolved mode for more information.	
<code>includeInstanceUpdates</code>	Defines if the content updates of the dataset are included. Default is <code>false</code> .	No
<code>pagination</code>	Enables pagination context for the operations <code>getChanges</code> and <code>getChangesOnDataSet</code> . 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. Get changes persisted context is built at first call according to the page size parameter in request. The pagination context is persisted on the server file system and allows invoking the next page until last page or when a timeout is reached. For creation: Defines <code>pageSize</code> parameter. For next: Defines <code>context</code> element with <code>identifier</code> from previous response. Enables pagination, see child elements below.	No
<code>pageSize</code> (nested under <code>pagination</code> element)	Defines maximum number of records in each page. Minimal size is <code>50</code> .	No (Only for creation)
<code>context</code> (nested under <code>pagination</code> element)	Defines content of pagination context.	No (Only for next)

ELEMENT	DESCRIPTION	REQUIRED
identifier (nested under <code>context</code> element)	Pagination context identifier.	Yes
<code>disableRedirectionToLastBroadcast</code>	See the description under Common parameters .	

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`

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> . 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 <code>updated</code> element)	Only the group of field have been updated.	Yes
change (nested under <code>changes</code> element)	Group of fields have been updated with own <code>XPath predicate</code> attribute of the record.	Yes
path (nested under <code>change</code> element)	Path in the record.	Yes
path (nested under <code>changes</code> element)	Path in the dataset.	Yes

ELEMENT	DESCRIPTION	REQUIRED
data (nested under <code>updated</code> 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. Content corresponding to a list of <code>predicate</code> element who contains the XPath predicate of record.	No
pagination	When pagination is enabled on request. Get changes persisted context allows invoking the next page until last page or when the context timeout is reached. Contains a next page: Defines <code>context</code> element with <code>identifier</code> . Is the last page: Defines <code>context</code> element without <code>identifier</code> . Enables pagination, see child elements below.	No
context (nested under <code>pagination</code> element)	Defines content of pagination context.	Yes (Only for next and last)
identifier (nested under <code>context</code> element)	Pagination context identifier. Not defined at last returned page.	No
pageNumber (nested under <code>context</code> element)	Current page number in pagination context.	Yes
totalPages (nested under <code>context</code> 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>

    <disableRedirectionToLastBroadcast>boolean</
disableRedirectionToLastBroadcast>

</m:getCredentials_{TableName}>

```

with:

ELEMENT	DESCRIPTION	REQUIRED
branch	See the description under Common parameters .	
version	See the description under Common parameters .	
instance	See the description under Common parameters .	
predicate	See the description under Common parameters .	
viewPublication	See the description under Common parameters .	

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](#).

```

<m:multi_ xmlns:m="urn:ebx-schemas:dataservices_1.0">

    <branch>String</branch>

    <version>String</version>

```

```

    <instance>String</instance>

    <blockingConstraintsDisabled>boolean</blockingConstraintsDisabled>

    <disableRedirectionToLastBroadcast>boolean</
disableRedirectionToLastBroadcast>

    <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 Common parameters .	
version	See the description under Common parameters .	

ELEMENT	DESCRIPTION	REQUIRED
instance	See the description under Common parameters .	
blockingConstraintsDisabled	See the description under Common parameters .	
details	See the description under Common parameters .	
disableRedirectionToLastBroadcast	See the description under Common parameters .	
request	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. Operations such as <code>count</code> , <code>select</code> , <code>getChanges</code> , <code>getCredentials</code> , <code>insert</code> , <code>delete</code> or <code>update</code> .	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`).

See also:

- [Procedure](#)
- [Repository](#)

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	This element contains the response of one operation. It is be repeated multiple times for additional operations. Each response is

ELEMENT	DESCRIPTION
	identified by an 'id' attribute set in the request or automatically generated. The content of the element corresponds to the response of a single operation, such as <code>count</code> , <code>select</code> , <code>getChanges</code> , <code>getCredentials</code> , <code>insert</code> , <code>delete</code> or <code>update</code> .

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 `includesMetadata` with `system` value:

```
<m:select_{TableName} xmlns:m="urn:ebx-schemas:dataservices_1.0">

  <branch>String</branch>

  <version>String</version>

  <instance>String</instance>

  <predicate>String</predicate>

  <includesMetadata>system</includesMetadata>

</m:select_{TableName}>
```

Use the value of the `ebx-metadata/system/update_time` element in the following update request. If the record was changed since the specified time, the update cancels. To prevent update of a modified record, add the `ebx-metadata/system/update_time` element value.

```
<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>

      <a>String</a>

      <b>String</b>

      <c>String</c>

      <d>String</d>

      ...

      <ebx-metadata>

        <system>

          ...

          <update_time>2010-06-28T10:10:31.046</update_time>

        </system>

      </ebx-metadata>

    </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:

You can use the `checkNotChangedSinceLastTime` element 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.

Operations on datasets and dataspace

Parameters for operations on dataspace and snapshots are as follows:

ELEMENT	DESCRIPTION	REQUIRED
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

<i>ELEMENT</i>	<i>DESCRIPTION</i>	<i>REQUIRED</i>
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	Defines if validation returns details. 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. The optional attribute <code>locale</code> (default 'en-US') defines the language in which the validation messages are to be returned.	No. If not specified, no details are returned.
owner	Defines the owner. Must respect the inner format as returned by <code>Profile.format</code> .	No

<i>ELEMENT</i>	<i>DESCRIPTION</i>	<i>REQUIRED</i>
branchToCopyPermissionFrom	Defines the identifier of the dataspace from which to copy the permissions.	No
documentation	Documentation for a dedicated language. Multiple <code>documentation</code> elements may be used for several languages.	No
locale (nested under the <code>documentation</code> element)	Locale of the documentation.	Only required when the <code>documentation</code> element is used
label (nested under the <code>documentation</code> element)	Label for the language.	No
description (nested under the <code>documentation</code> 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>

                <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 <code>message</code> elements may be used.	No
locale (nested under the <code>message</code> element)	Locale of the user message.	Only required when the <code>message</code> element is used

<i>ELEMENT</i>	<i>DESCRIPTION</i>	<i>REQUIRED</i>
label (nested under the <code>message</code> 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:

<i>ELEMENT</i>	<i>DESCRIPTION</i>
status	'00' indicates that the operation has been executed successfully. '94' indicates that the dataspace has been already locked by another user. 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 dataspace 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. Otherwise, a SOAP exception is thrown.

Merge a dataspace

Merge dataspace 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 dataspace operation. Sets whether the specified dataspace and its associated snapshots will be deleted upon merge. When this parameter is not specified in the request, the default value is <code>false</code>. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.dataDeletionOnCloseOrMerge.default</code> in the EBX® main configuration file. See	No

ELEMENT	DESCRIPTION	REQUIRED
	Deleting data and history for more information.	
deleteHistoryOnMerge	This parameter is available for the merge dataspace operation. Sets whether the history associated with the specified dataspace will be deleted upon merge. Default value is <code>false</code>. When this parameter is not specified in the request, the default value is <code>false</code>. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.historyDeletionOnCloseOrMerge.default</code> in the EBX® main configuration file. See Deleting data and history 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 dataspace automatically overrides the data in the parent dataspace.

Merge dataspace 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	This parameter is available for the close dataspace and close snapshot operations. Sets whether the specified snapshot, or dataspace and its associated snapshots, will	No

ELEMENT	DESCRIPTION	REQUIRED
	be deleted upon closure. When this parameter is not specified in the request, the default value is <code>false</code>. It is possible to redefine this default value by specifying the property <code>ebx.dataservices.dataDeletionOnCloseOrMerge.default</code> in the EBX® main configuration file. See Deleting data and history for more information.	
deleteHistoryOnClose	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 <code>false</code>. When this parameter is not specified in the request, the default value is <code>false</code>. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.historyDeletionOnCloseOrMerge.default</code> in the EBX® main configuration file. See Deleting data and history for more information.	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:

<i>ELEMENT</i>	<i>DESCRIPTION</i>	<i>REQUIRED</i>
branch	See the description under Common parameters .	Yes
instance	See the description under Common parameters .	Yes
unitName	Name of the replication unit. See also:	Yes

ELEMENT	DESCRIPTION	REQUIRED
	Replication refresh information	

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.

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](#) for more information.

ELEMENT	DESCRIPTION	REQUIRED
parameters	<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. Note:	No

ELEMENT	DESCRIPTION	REQUIRED
	The <code>parameters</code> element is ignored if at least one session parameter has been defined.	
parameter (nested under the <code>parameters</code> element). Multiple <code>parameter</code> elements may be used.	An input parameter for the workflow.	No
name (nested under the <code>parameter</code> element)	Name of the parameter.	Yes
value (nested under the <code>parameter</code> 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 Common parameters .	No
parameters	<i>Deprecated since version 5.7.0</i> See the description under Common parameters .	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	No

ELEMENT	DESCRIPTION	REQUIRED
	compatibility, it will eventually be removed in a future major version.	
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
resumeld	Resume identifier of the waiting task.	Yes
parameters	<i>Deprecated since version 5.7.0</i> See the description under Common parameters .	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. '20' indicates that the workflow has not been found. '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.

ELEMENT	DESCRIPTION	REQUIRED
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

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](#) 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](#) for more information.

User interface operations

See [Application locking](#) 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
<code>details</code>	Defines attributes that must be applied to response messages. The attribute <code>locale</code> (default: EBX® default <code>locale</code>) defines the language in which the system	No, but if specified, the <code>locale</code> attribute must be provided.

PARAMETER	DESCRIPTION	REQUIRED
	item messages must be returned.	

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

- [Introduction](#)
- [Built-in RESTful services](#)
- [JSON Formats](#)
 - [Introduction](#)
 - [Extended](#)
 - [Compact](#)
 - [Common](#)
 - [Others](#)

Introduction

- [Overview](#)
- [Activation and configuration](#)
- [Interactions](#)
 - [Body media types](#)
 - [Tracking information](#)
 - [Session parameters](#)
 - [Session channel](#)
 - [Procedure context](#)
 - [Exception handling](#)
- [Security](#)
 - [Authentication](#)
 - [Global permissions](#)
 - [Lookup mechanism](#)
- [Monitoring](#)
- [SOAP and REST comparative](#)

- Limitations
 - Date, time & dateTime format
 - JMS

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](#).

Built-in REST data services allow performing operations such as:

- Selecting, inserting, updating, deleting, or counting records
- Selecting or counting history records
- Selecting, updating, or counting dataset values
- Selecting or updating dataspace or snapshot information
- Selecting children dataspaces or snapshots
- Creating, merging, or closing a dataspace
- Creating or closing a snapshot
- Administrative operations to manage access to the UI or to system information

Note:

See [SOAP and REST comparative](#).

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](#) for more information.

In case of specific deployment, for example using reverse-proxy mode, see [URLs computing](#) for more information.

Only HTTP(S) protocols are supported.

Interactions

Body media types

Input and output messages must be *exclusively* in JSON format with UTF-8 encoding for REST built-in. However, input message can also be used to pass HTML form data.

See [Content-Type](#) for more information.

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 extended JSON format request contains:

```
{
    "procedureContext":           // JSON Object (optional)
    {
        "trackingInformation": "String" // JSON String (optional)
    },
    ...
}
```

For more information, see `Session.getTrackingInfo` 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 extended JSON format 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.getInputParameterValue` in the Java API.

Session channel

The session channel allows to filter what can be selected or modified, from the EBX® repository, when using a REST built-in or REST toolkit service. The filter is based on table, group or field configuration where the visibility is defined through the data model, by specifying a [default view](#).

It can be specified through the query parameter `ebx-channel`. It's available values are:

- `dataServices`

- `ui`

The filter behavior is described by this combinatorial:

Data services channel

XML ELEMENT	VALUE	SCHEMA NODE TYPE	VIEW	BEHAVIOUR
<code><hiddenInDataset true</code>	<code>true</code>	<code>field node</code>	Default tabular view	Hidden for content and not sortable
CustomView (tabular or hierarchical)	Hidden for meta, content, filter and sort			
Default form record	Hidden for meta and content			
Default form record field	Not found			

User Interface channel

XML ELEMENT	VALUE	SCHEMA NODE TYPE	VIEW	BEHAVIOUR
<code><hidden></code>	true	table node	Tree view	hidden for meta and content
Default tabular view	not found			
CustomView (tabular or hierarchical)	forbidden			
Default form record	not found			
Default form record field				

XML ELEMENT	VALUE	SCHEMA NODE TYPE	VIEW	BEHAVIOUR
Custom form record				
field node	Default tabular view	hidden for content and not sortable		
Default form record	hidden for content			
Default form record field	not found			
<hiddenInViews>	true	field node	CustomView (tabular or hierarchical)	hidden for meta, content, filter and sort
Custom form record	hidden for meta and content			
<hiddenInSearch>	true	field node	Default tabular view	not filterable
CustomView (tabular or hierarchical)				
Default form record				
Default form record field				
Custom form record				
textSearchOnly	Default tabular view	not filterable except for text search		

XML ELEMENT	VALUE	SCHEMA NODE TYPE	VIEW	BEHAVIOUR
CustomView (tabular or hierarchical)				
Default form record				
Default form record field				
Custom form record				

Note:

The above field nodes can only be under table nodes.

Procedure context

Depending on the data services operation being called, it is possible to overwrite the default procedure context configuration. They are defined in the request body and are applied within the built-in operation.

Procedure context can be applied to custom REST Toolkit services.

- Example for a RESTful operation, the JSON request body contains:

```
{
  "procedureContext":           // JSON Object (optional)
  {
    "commitThreshold": Integer  // JSON Number (optional)
  },
  ...
}
```

```
}
```

For more information, see `ProcedureContext.setCommitThreshold`, `SessionContext.getProcedureUtility` and `ProcedureUtility.execute` in the Java API.

Exception handling

When an error occurs, a JSON exception response is returned to the caller. In most cases the exception takes JSON format like in the REST Toolkit.

See also:

- [REST Toolkit exception handling](#)

```
{
  "errorCode": 999,          // JSON Number, EBX® error code used mainly
  for                        // HTTP error 422 (optional)

  "localizedMessage": "...", // JSON String, localized message

  "messages": [              // JSON Array of String, Internal messages
    useful when debugging    // (optional).

                                // Usually not displayed to the user and not
    localized.

    "Message 1", "Message 2"

  ]
}
```

If there are errors when [inserting or updating multiple records](#), an error report is returned.

```

{
  "code": 999,          // JSON Number, HTTP status code

  "errors": [

    {

      "level": "...",    // JSON String, severity level (optional)

      "rowIndex": 999,    // JSON Number, request row index (optional)

      "userCode": "...",  // JSON String, user code (optional)

      "message": "...",   // JSON String, message

      "blocksCommit": "...", // JSON String, Type of blocking constraints
(optional)

      "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 level of each error is specified by a character, with one of the possible values (`fatal` , `error` , `warning` , `info`).

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
- `Severity.getLabel`

Security

Authentication

Authentication is mandatory to access to the content. 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](#).

See also:

- [HTTP Authorization header policy](#)

- 'Basic Authentication Scheme' method is based on the HTTP-Header `Authorization` in base64 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 QWxpYmFiYTpvYVUHNlc2FtZQ==
```

Note:

The `WWW-Authenticate` header can be valued with this method.

See also:

- [HTTP Authorization header policy](#)

- '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](#) section.

For more information on this authentication scheme, see

`Directory.authenticateUserFromLoginPassword`.

- The 'REST Forward Authentication Scheme' is used only when calling a REST service from a [user service](#), that reuses the current authenticated session.

For more information, see [Implementing a user service](#) making a call to [REST data services](#).

- '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` for more information.
- 'Anonymous authentication Scheme' is used only to access the REST services handling the authentication operations. The credentials acquisition, password changes, etc. implies that the user cannot be known yet.

Global permissions

Global access permissions can be independently defined for the REST built-in and REST OpenAPI services. See [Global permissions](#) for more information.

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	Token • The HTTP request must hold an <code>Authorization</code> header. • <code>Authorization</code> header value must start with the word <code>EBX</code>. • No <code>login</code> is provided in the URL parameters. Basic • The HTTP request must hold an <code>Authorization</code> header. • <code>Authorization</code> header value must start with the word <code>Basic</code>. • No <code>login</code> is provided in the URL parameters. Standard • The HTTP request must not hold an <code>Authorization</code> header. • A <code>login</code> and a <code>password</code> are provided in the URL parameters. Rest forward • The HTTP request must not contain an <code>Authorization</code> header. • No <code>login</code> is provided in the URL parameters. Specific

OPERATION / PROTOCOL	AUTHENTICATION METHODS AND APPLICATION CONDITIONS
	- The HTTP request must not satisfy the conditions of the previous authentication methods. Anonymous - None of the previous authentication methods can be applied. - The requested REST service is handling an authentication operation.

In case of multiple authentication methods present in the same request, EBX® will return an HTTP code `401 Unauthorized`.

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](#)
- [TIBCO EBX® main configuration file](#)

SOAP and REST comparative

OPERATIONS	SOAP	REST
Data		
Select metamodel		X
Select or count records (with filter and/or view publication)	X	X

OPERATIONS	SOAP	REST
Select records with association values (first level)	X	
Update or delete records (with filter and/or view publication)		X
Selector for possible enumeration values (with filter)		X
Prepare for creation or duplication		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	
Generate service contract	WSDL	OpenAPI
Views		
Look up published table views		X

OPERATIONS	SOAP	REST
Dataspaces		
Select dataspace or snapshot information		X
Select root or children dataspaces, or select snapshots		X
Create, close, merge a dataspace	X	X
Create, close a snapshot	X	X
Validate a dataspace or a snapshot	X	
Validate a dataset	X	
Locking, unlocking a dataspace	X	X
Workflow		
Start, resume or end a workflow	X	
Administration		
Manage the default directory content 'Users', 'Roles'... tables.	X	X
Open, close the user interface	X	X
Select, insert, update, delete operations for administration dataset		X
Select the system information	X	X

OPERATIONS	SOAP	REST
Staging API		X (*)
Other		
Develop web services from the Java API		X (*)

(*) See [REST Toolkit](#) for more information.

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.

Built-in RESTful services

- [Introduction](#)
- [Request](#)
 - [HTTP method](#)
 - [URL](#)
 - [Header fields](#)
 - [Common parameters](#)
 - [Message body](#)

- Response
 - Header fields
 - HTTP codes
 - Message body
- Administration operations
 - Directory operations
 - User interface operations
 - System information operation
- Token authentication operations
 - Create token operation
 - Change password operation
 - Revoke token operation
- Data operations
 - Select operation
 - Prepare operations
 - Insert operation
 - Update operation
 - Delete operation
 - Count operation
 - Optimistic locking
 - Inheritance
 - Look up table views operation
- Form data operations
 - Form insert operation
 - Form update operation
- Beta feature: Dataspace operations
 - Beta feature: Select dataspace or snapshots

- Beta feature: Create a child dataspace or a snapshot
- Beta feature: Locking a dataspace
- Beta feature: Unlocking a dataspace
- Beta feature: Merge a dataspace
- Beta feature: Close a dataspace or a snapshot
- Health operations
 - Overview
 - Started operation
 - Check operation
 - Stopped operation
- OpenAPI operations
 - Overview
 - Generate OpenAPI document
- Staging operations
 - Domain management operations
 - Staging element navigation operations
 - Archive manipulation operations
 - Export operations
 - Import operations
- Limitations
 - General limitations
 - Compact format limitations
 - Read operations
 - Write operations
 - Directory operations
 - OpenAPI operations

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 for an intuitive and straightforward navigation, which implies that the data details could be obtained through a link.

Note:

All operations are stateless.

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](#)
- [Security](#)

HTTP method

Considered HTTP methods for built-in RESTful services, are:

- **GET** : used to select master data defined in the URL (the size limit of the URL depends on the application server and on the browser; it must be lower than or equal to 2kB).
- **POST** : used to insert 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 by the table URL and the record keys 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](#).
- `{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](#) or dedicated operation parameters passed in the URL.

Operation category

Specializes the operation; it is added in the path of the URL in `{category}`, and takes one of the following values:

<code>admin</code>	Administration operations reserved to administrators. See Administration operations for more information.
<code>auth</code>	Manages token authentication method. See Token authentication operations and Token Authentication Scheme for more information.
<code>data</code> or <code>data-compact</code>	Lists dataset content; requests a table, a record or a field record content, including modified operations on dataset node, table,

	record and record field.(A compact format is available, to ease interaction in simple use cases.) Manages dataspace and snapshot life cycle. See Data operations, Beta feature: Dataspace operations and Compact format limitations for more information.
<code>form-data</code> or <code>form-data-compact</code>	Validates incoming data, and returns a report before inserting or updating a dataset node, record or record field. (A compact format is available to ease interaction for simple use cases.) See Form data operations and Compact format limitations for more information.
<code>health</code>	Server health information returns information for a monitoring tool or container probes. See Health operations for more information.
<code>history</code>	Lists the content of a history dataset; requests a history table, a history of a record or a history record. See Data operations for more information. See also: • History
<code>api</code>	Generates the OpenAPI document of the selected resource. See OpenAPI operations for more information.

Header fields

These header field definitions are used by TIBCO EBX®.

Accept	Used to specify (by order of preference) content types to be used in the response: the first supported one will be chosen and specified in the response header <code>Content-Type</code>. Currently, the only supported one is <code>application/json</code>. If none is supported, the result depends on property <code>ebx.dataservices.rest.request.checkAccept</code>: • If <code>true</code>, an HTTP error response is returned with code <code>406</code>. • If <code>false</code>, the response is returned with the default content type, that is <code>application/json</code>. See also: • Configuring data services
Accept - Language	Used to specify the preferred locale for the response. The supported locales are defined in the schema model. If none of the preferred locales is supported, the default locale for the current model is used.
Authorization	Supported authentication schemes include 'Basic Authentication Scheme' and 'Token Authentication Scheme'; if another scheme is used, the request is rejected. See also:

	• Authentication
<code>Content-Type</code>	Used to specify the request body media type. The supported types are <code>application/json</code> and <code>application/x-www-form-urlencoded</code>. If the request value is not supported, then an HTTP error message is returned with the code <code>415</code> (<i>Unsupported media type</i>). See also: • Configuring data services
<code>X-Requested-With</code>	If present and in case of authentication failure, prevents from adding the <code>WWW-Authenticate</code> header in the response. See also: • Response header <code>WWW-Authenticate</code>

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>	This parameter only has impact in a D3 architecture. 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.

PARAMETER	DESCRIPTION
	If the specified dataspace is not a delivery dataspace on a D3 primary node, this parameter is ignored. <code>Boolean</code> type value. If this parameter is not present, the default is <code>false</code> (redirection to a D3 master enabled), unless the configuration property <code>ebx.dataservices.disableRedirectionToLastBroadcast.default</code> has been set. See also: • Primary node
<code>ebx-indent</code> <code>indent</code> (<i>Deprecated since 6.0.0</i>)	Specifies whether the response should be indented, to be easier to read for a human. <code>Boolean</code> type, default value is <code>false</code>.
<code>ebx-channel</code>	Specifies the session channel. <code>String</code> type, possible values are: • <code>dataServices</code> • <code>ui</code> The default value is <code>dataServices</code>. See also: • <code>Constants.Data.PARAM_EBX_CHANNEL</code>

Message body

It contains the request data using the JSON format, see [Extended JSON request body](#) and [Compact JSON request body](#).

Note:

Requests may define a message body only when using `POST` or `PUT` HTTP methods.

Response

This chapter describes the responses returned by built-in RESTful services.

- See [Exception handling](#) 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®.

<code>Content - Language</code>	Indicates the locale used in the response for labels and descriptions.
<code>Content - Type</code>	Indicates the response body content type.
<code>Location</code>	If a new record has been successfully inserted, the query URL for this record is returned by this field.
<code>WWW-Authenticate</code>	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: • the 'Basic Authentication Scheme' method is enabled and • the <code>X-Requested-With</code> HTTP header is not present.

If the client is able to interpret the authentication method, it is possible to resubmit the request providing the appropriate credentials.

The administration property `ebx.dataservices.rest.auth.tryBasicAuthentication` must be set to `true`.

See also:

- [Request header](#) `X-Requested-With`
- [Authentication](#)

HTTP codes

HTTP CODE	DESCRIPTION
<code>200</code> (<i>OK</i>)	The request has been successfully handled.
<code>201</code> (<i>Created</i>)	A new record has been created, in this case, the header field <code>Location</code> is returned with its resource URL.
<code>204</code> (<i>No content</i>)	The request has been successfully handled but no response body is returned.
<code>400</code> (<i>Bad request</i>)	The request URL or body is not well-formed or contains invalid content.
<code>401</code> (<i>Unauthorized</i>)	Authentication has failed.
<code>403</code> (<i>Forbidden</i>)	Permission was denied to read or modify the specified resource for the authenticated user. This error is also returned when the user: • is not allowed to modify a field mentioned in the request message body.

HTTP CODE	DESCRIPTION
	is not allowed to access the REST connector. For more details, see Global permissions.
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>	A concurrent modification has occurred. See also: Optimistic locking
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 ([Extended JSON](#) and [Compact JSON](#) samples).

With the exception of code `204` (*No content*).

- HTTP codes greater than or equal to `300`: the content describes the error. See [JSON](#) for details on the format.

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](#).

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 metamodel is requested.

See [select](#) and [update](#) 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](#) and [update](#) 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](#), [update](#), [insert](#) and [delete](#) 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 metamodel is requested.

See [select](#), [update](#), [insert](#) and [delete](#) 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'.

The URL format is:


```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/Bebx-manager/ebx-  
manager/domain/toolStatus
```

See also:

- [Application locking](#)

Retrieve user interface state

User interface status and the unavailability message are accessible like dataset nodes.

See [Select operation](#) and the [Extended JSON](#) or [Compact JSON](#) 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](#) and the [Extended JSON](#) or [Compact JSON](#) 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
- Repository administration

Parameters

The following parameter is applicable.

PARAMETER	DESCRIPTION
<code>systemInformationMode</code>	Specifies the returned mode: • <code>flat</code> : A flat representation under the following information groups: <code>bootInfoEBX</code> , <code>repositoryInfo</code> and <code>bootInfoVM</code> . • <code>hierarchical</code> : A hierarchical representation. <code>String</code> type, default value is <code>flat</code> .

HTTP codes

HTTP CODE	DESCRIPTION
<code>200</code> (<i>OK</i>)	The system information was successfully returned.
<code>400</code> (<i>Bad request</i>)	The request is not correct, it contains one of the following errors: • the HTTP method is not <code>GET</code> nor <code>POST</code> , • the HTTP parameter <code>systemInformationMode</code> is not correct, • the operation is not supported.

HTTP CODE	DESCRIPTION
	the request path is invalid.
403 (Forbidden)	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](#) example of the flat representation.

See the [JSON](#) example of the hierarchical representation.

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` (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](#).

The 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.authenticateUserFromLoginPassword` 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.authenticateUserFromHttpRequest` for more information.

See the [JSON](#) examples of a token creation request.

HTTP codes

HTTP CODE	DESCRIPTION
<code>200</code> (<i>OK</i>)	The token was successfully created.
<code>400</code> (<i>Bad request</i>)	For one of the following reasons: • the syntax is not correct, • the HTTP method is not <code>POST</code>, • the operation is not supported.
<code>401</code> (<i>Unauthorized</i>)	For one of the following reasons: • The <code>login</code> and/or <code>password</code> is/are incorrect. • Authentication data for other methods is defined.
<code>422</code> (<i>Unprocessable entity</i>)	For one of the following reasons: • <code>PasswordMustChange</code>: The password must be changed (only available with the default directory). See Change password operation.

HTTP CODE	DESCRIPTION
	<code>RestrictedAccess</code>: User access is closed for maintenance or other actions (reserved to administrators).

Response body

If the HTTP code is `200 (OK)`, the body holds the token value and its type.

See the [JSON](#) 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

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](#) 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

The 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](#) example of a password change and token creation request.

HTTP codes

HTTP CODE	DESCRIPTION
204 (No content)	The password has been changed.
400 (Bad request)	For one of the following reasons: - the EBX® default directory is required, - the syntax is not correct, - the HTTP method is not <code>POST</code>, - the provided <code>login</code> and the user's session one mismatch - the operation is not supported.
401 (Unauthorized)	For the following reason: The <code>login</code> and/or <code>password</code> is/are incorrect.
422 (Unprocessable entity)	For one of the following reasons: <code>PasswordChangeAbort</code> : <code>passwordNew</code> is empty. <code>PasswordChangeAbort</code> : <code>passwordNew</code> is equal to <code>password</code>. <code>PasswordChangeAbort</code> : <code>password</code> is incorrect.

HTTP CODE	DESCRIPTION
	<code>RestrictedAccess</code>: User access is closed for maintenance or other actions (reserved to administrators).

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.

The URL format is:

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/auth/v1/token:revoke
```

Header fields

<code>Authorization</code>	This field is required, <code>tokenType</code> and <code>accessToken</code> fields must have the values returned from the "token create" operation. > Authorization: <tokenType> <accessToken>

HTTP codes

HTTP CODE	DESCRIPTION
<code>204 (No content)</code>	The token has been revoked successfully.
<code>400 (Bad request)</code>	For one of the following reasons: - the configuration is not activated, - the syntax is incorrect, - the HTTP method is not <code>POST</code>,

HTTP CODE	DESCRIPTION
	the operation is not supported.
401 (<i>Unauthorized</i>)	Authentication has failed.

Data operations

The `data` category operations concern the datasets, the dataset fields, tables, records or record fields.

The `data-compact` category operations concern the dataset fields, tables, records or record fields.

The `history` category operations concern historized content from datasets, tables, records or record fields.

The `form-data` category operations concern the dataset fields, records or record fields when constraints must remain valid on content creation or update.

The `form-data-compact` category operations concern the dataset fields, records or record fields when constraints must remain valid on content creation or update.

See [Form data operations](#) for more information.

Select operation

Select operation returns hierarchical content. This operation may use one of the following methods:

- `GET` HTTP method,
- `POST` HTTP method without message body or
- `POST` HTTP method with a message body, with `:select` URL extended action and optionally [session parameters](#).

URL formats are:

- Dataset tree**, depending on [operation category](#):

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}[:select]
```

Note:

Terminal nodes and sub-nodes are not included.

- **Dataset node:** the **data** or **data-compact** category returns the terminal nodes contained in the selected node.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}[:select]
```

Note:

Not applicable with the **history** category.

- **Table**, depending on **operation category**:

the **data** or **data-compact** category returns the table content and/or metamodel, current page records and URLs for pagination.

The **history** category returns the history table content and/or metamodel, current page records and URLs for pagination.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}[:select]
```

See also:

- Count operation
- Look up table views operation

- **Record**, depending on [operation category](#):

the `data` or `data-compact` category returns the record content and/or metamodel.

The `history` category returns history record content and/or metamodel.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}[:select]
```

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}[:select]?
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](#):

the `data` or `data-compact` 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}
[:select]
```

Note:

The field must be either an association node, a selection node, a terminal node or above.

Where:

- `{category}` corresponds to the [operation category](#) (possible values are: `data` or `data-compact`).
- `{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 percent encoded representation of the primary key (see [RFC-3986 Uniform Resource Identifier](#)).
- `{xpathExpression}` corresponds to the record primary key, using the XPath expression.
- `{pathInRecord}` corresponds to the path starting from the table node.
- `:select` extended action is required when POST HTTP method is used with a body message.

Parameters

The following parameters are applicable to the select operation.

PARAMETER	DESCRIPTION
<code>includeContent</code>	Includes the <code>content</code> field with the content corresponding to the selection. <code>Boolean</code> type, default value is <code>true</code> .
<code>includeDetails</code>	Includes the <code>details</code> field in the metamodel and in the content, for each indirect reachable resource. The returned value corresponds to its URL resource.

PARAMETER	DESCRIPTION
	Type <code>Boolean</code>, default value is <code>true</code>. See also: • includeMetamodel
<code>includeHistory</code>	Includes those fields for a historized content: • The <code>history</code> property in the metamodel. The returned value corresponds to a boolean value. • The <code>historyDetails</code> property in the content and for each indirectly reachable resource. This point is coupled to the includeDetails parameter. The returned value corresponds to its URL resource. <code>Boolean</code> type, default value is <code>false</code>. Note: The <code>includeHistory</code> parameter is ignored in the <code>history</code> category, the default value is <code>true</code>. See also: • includeMetamodel • includeDetails
<code>includeLabel</code>	Includes the <code>label</code> field associated with each simple type <code>content</code>. Possible values are: • <code>yes</code>: the <code>label</code> is included for the foreign key, enumeration, record and selector values.

PARAMETER	DESCRIPTION
	• <code>all</code> : the <code>label</code> field is included, as for the <code>yes</code> value and also for the Content of simple type. See also: ◦ <code>SchemaNode.displayOccurrence</code> • <code>no</code> : the <code>label</code> field is not included (integration use case). <code>String</code> type, default value is <code>yes</code>. Note: The <code>label</code> field is not included if it is equal to the <code>content</code> field.
<code>includeMeta</code>	<i>Deprecated since 6.1.0, replaced by includeMetamodel.</i>
<code>includeMetadata</code>	The response contains the specified metadata. Each returned record contains additional elements called steps under the <code>ebx-metadata</code> root element. For instance, the 'system' step value will contain technical data. See optimistic locking for more information. String value (default value is empty) steps must be separated by a comma: <code>system, teamUp</code>. All steps can be returned using the <code>ebx-all</code> value. See also: • Technical data

PARAMETER	DESCRIPTION
<code>includeMetamodel</code>	Includes the <code>meta</code> field corresponding to the description of the structure returned in the <code>content</code> field. <code>Boolean</code> type, default value is <code>false</code>. See also: • Metamodel • includeHistory • includeDetails
<code>includeMergeInfo</code>	Includes the <code>merge_info</code> corresponding to a field of the technical data of an history transaction and has a potentially high access cost. <code>Boolean</code> type, default value is <code>true</code>. Note: This parameter is ignored with the <code>data</code> category. See also: • REST access to history table
<code>includeOpenApiDetails</code>	Includes the OpenAPI specification URL for each describable node. Type <code>Boolean</code>, default value is <code>false</code>. Note: This query parameter is ignored if the includeDetails one is set to <code>false</code>. See also:

PARAMETER	DESCRIPTION
	• OpenAPI operations • includeDetails
<code>includeSelector</code>	Includes the <code>selector</code> field in the response, for each indirect reachable resource. The returned value corresponds to its URL resource. Type <code>Boolean</code>, default value is <code>true</code>. See also: • selector
<code>includeSortCriteria</code>	Includes the <code>sortCriteria</code> field corresponding to the list of sort criteria applied. The sort criteria parameters are added by using: • sort • sortOnLabel • viewPublication <code>Boolean</code> type, default value is <code>false</code>. Example JSON
<code>includeTechnicals</code>	<i>Deprecated since 6.1.0, replaced by <code>includeMetadata</code>.</i> Note: This parameter is ignored with the <code>history</code> category.

PARAMETER	DESCRIPTION
<code>includeValidation</code>	Includes the validation report corresponding to the selection. <code>Boolean</code> type, default value is <code>false</code>. Note: This parameter is ignored with the <code>history</code> category. See also: includeDetails

Table parameters

The following parameters are applicable to tables, associations and selection nodes.

PARAMETER	DESCRIPTION
<code>filter</code>	Quick search predicate or complete XPath predicate expression defines the field values to which the request is applied. If empty, all records will be retrieved. <code>String</code> type value. Note: The history code operation value is usable with <code>ebx-operationCode</code> path field from the <code>meta</code> section associated with this field. See also: sort by relevancy <code>includeMetadata</code> parameter See also:

PARAMETER	DESCRIPTION
	• Search predicate
<code>historyMode</code>	Specifies the filter context applied on table. <code>String</code> type, possible values are: • <code>CurrentDataSpaceOnly</code> : history in current dataspace • <code>CurrentDataSpaceAndAncestors</code> : history in current dataspace and ancestors • <code>CurrentDataSpaceAndMergedChildren</code> : history in current dataspace and merged children • <code>AllDataSpaces</code> : history in all dataspaces The default value is <code>CurrentDataSpaceOnly</code>. See also: • <code>history</code> Note: This parameter is ignored with the <code>data</code> category.
<code>includeOcculting</code>	Includes the records in occulting mode. <code>Boolean</code> type, default value is <code>false</code>. See also: • Inheritance

PARAMETER	DESCRIPTION
<code>primaryKey</code>	Search a record by a primary key, using the XPath expression. The XPath predicate expression should only contain field(s) of the primary key and all of them. Fields are separated by the operator <code>and</code>. A field is represented by one of the following possibilities according to its simple type: - For the <code>date</code>, <code>time</code> or <code>dateTime</code> types: use the <code>date-equal(path, value)</code> - For other types: indicate the <code>path</code>, the <code>=</code> operator and the <code>value</code>. Example with a composed primary key: <code>./pk1i=1 and date-equal(./pk2d, '2015-11-13')</code> The response will only contain the corresponding record, otherwise an error is returned. Consequently, the other table parameters are ignored (as filter, viewPublication, sort, etc.) <code>String</code> type value.
<code>pageFirstRecordFilter</code>	<i>Deprecated since 5.9.0, replaced by <code>pageRecordFilter</code></i>
<code>pageRecordFilter</code>	Specifies a server side built filter, for pagination, pointing to a record. This filter is strongly linked to the <code>pageAction</code> value and should not be modified on the client side. The filter takes the form of a record XPath predicate expression used to figure out the pagination contexts.

PARAMETER	DESCRIPTION
	<code>String</code> type. See also: • Pagination
<code>pageAction</code>	Specifies the pagination action to perform from the identifier held by <code>pageRecordFilter</code>. <code>String</code> type, default value is <code>first</code>. The possible values are: • <code>first</code> • <code>previous</code> • <code>next</code> • <code>last</code> See also: • <code>RequestPagination</code>
<code>pageSize</code>	Specifies the maximum number of records per page. <code>Integer</code> type, default value is based on the user preferences. If the specified value is <code>0</code>, selects with the maximum recommended page size, default is 10000. See also: • ebx.dataservices.pagination.maxSize.default Note:

PARAMETER	DESCRIPTION
	If the specified page size value exceeds the maximum page size, the maximum page size is selected.
sort	Specifies that the operation result will be sorted according to the specified criteria. The criteria are composed of one or more criterion, 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: 1. the sortOnLabel parameter as a new criteria added after the sort . 2. the viewPublication parameter as a new criteria added after the sort . The value structure is as follows: <path1>:<order>;...;<pathN>:<order> Where: • <path1> corresponds to the field path in priority 1 . • <order> corresponds to the sorting order, with one of the following values: ◦ asc : ascending order on value (default), ◦ desc : descending order on value, ◦ lasc : ascending order on label, ◦ ldesc : descending order on label.

PARAMETER	DESCRIPTION
	<code>String</code> type, the default value orders according to the primary key fields (ascending order on value). Note: The history code operation value is usable with the <code>ebx-operationCode</code> path field from the <code>meta</code> section associated with this field. Note: This parameter is ignored when the sort by relevancy is activated. Note: <code>includeMetadata</code> parameter See also: <code>Request.setSortCriteria</code>
<code>sortByRelevancy</code>	Specifies that the operation result will be sorted by relevancy, only if these conditions are fulfilled: - The Quick Search is activated by specifying an <code>osd:search</code> predicate directly in the filter parameter. - The request is not applied on an inherited dataset, history table or mapped table. <code>String</code> type. The possible values are: <code>lasc</code> : ascending order on label, <code>ldesc</code> : descending order on label.

PARAMETER	DESCRIPTION
	If the sort by relevancy is activated, the following parameters are ignored: sort, sortOnLabel, sortPriority and sort criteria defined through the viewPublication. See also: • Search predicate
<code>sortOnLabel</code>	Specifies that the operation result will be sorted according to the record label. This parameter can be combined with: 1. the sort parameter as a new criteria added before the <code>sortOnLabel</code>. 2. the viewPublication parameter as a new criteria added after the <code>sortOnLabel</code>. The value structure is as follows: <code><order></code> Where: • <code><order></code> corresponds to the sorting order, with one of the following values: ◦ <code>1asc</code> : ascending order on label, ◦ <code>1desc</code> : descending order on label. The behavior of this parameter is described in the section defaultLabel. <code>String</code> type value. See also: • Limitation Note:

PARAMETER	DESCRIPTION
	This parameter is ignored when the sort by relevancy is activated.
<code>sortPriority</code>	Overrides the default priority of sort groups: • <code>sort</code> • <code>sortOnLabel</code> • <code>sortFromView</code> Comma separated <code>String</code> type, default value is: <code>sort,sortOnLabel,sortFromView</code>. Note: This parameter is ignored when the sort by relevancy is activated.
<code>viewPublication</code>	Specifies the name of the published view. This parameter can be combined with: 1. the filter parameter as the logical <code>and</code> operation. 2. the sort parameter as a new criteria added before the <code>viewPublication</code>. 3. the sortOnLabel parameter as new criteria added before the <code>viewPublication</code>. The behavior of this parameter is described in the section EBX® as a Web Component. <code>String</code> type value. Note:

PARAMETER	DESCRIPTION
	The sort criteria, defined through this parameter, is ignored when the sort by relevancy is activated. See also: • Look up table views operation

Selector parameters

The following parameters are only applicable to fields that return an enumeration, foreign key or `osd:resource` (Example [JSON](#)). By default, a pagination mechanism is always enabled. Some selector's select operations require input values. Thus, the uses of `POST` HTTP method [message body](#) allows to provide a record content.

PARAMETER	DESCRIPTION
<code>selector</code>	Specifies whether: • <code>true</code> : returns all possible values (includes their labels) • <code>false</code> : returns the current values for the current field. <code>Boolean</code> type, default value is <code>false</code>. Note: This parameter is ignored with the <code>history</code> category.
<code>firstElementIndex</code>	Specifies the index of the first element returned by the selector. Must be an integer higher than or equal to <code>0</code>. <code>Integer</code> type, default value is <code>0</code>.

PARAMETER	DESCRIPTION
<code>pageSize</code>	Specifies the maximum number of elements per page. <code>Integer</code> type, default value is based on the user preferences. If the specified value is <code>0</code>, selects with the maximum recommended page size, default is 10000. See also: ebx.dataservices.pagination.maxSize.default Note: If the specified page size value exceeds the maximum page size, the maximum page size is selected.
<code>selectorFilter</code>	Specifies the filter of the selector. <code>String</code> type value, the syntax complies with the Quick Search.

HTTP codes

HTTP CODE	DESCRIPTION
<code>200</code> (<i>OK</i>)	The selected resource is successfully retrieved.
<code>400</code> (<i>Bad request</i>)	The request is incorrect. This occurs when: - The selected field in a record or a dataset is sub-terminal, - The XPath predicate of the <code>filter</code> parameter is malformed or contains unfilterable nodes.

HTTP CODE	DESCRIPTION
	- The XPath predicate of the <code>primaryKey</code> parameter is malformed or is not a record primary key. - The sort criteria of the <code>sort</code> parameter have an invalid syntax or contain unsortable nodes. <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. <code>pageSize</code> parameter value is below 2. - The table view for the <code>viewPublication</code> parameter is either hierarchical, non-existent or non-published. - 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.
<code>403</code> (<i>Forbidden</i>)	The selected resource is hidden for the authenticated user.
<code>404</code> (<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.

Examples: [Extended JSON](#), [Compact JSON](#).

Prepare operations

Prepare for creation or duplication operations are used to create a new transient record with an initial content or with the content of a record to duplicate. A transient record corresponds to a content that is not persisted yet. Default values and fields, initialized by table triggers, are considered. Only field values with, at least, read-only access permissions are returned. These operations can optionally return metamodel to improve and assist data capture on the client side. Auto-incremented fields are only returned in the metamodel. By enabling the selector parameter, a transient record's field can be queried to the retrieve possible values of an enumerated field, a foreign key, etc. Furthermore, selector's select operations allows to provide input [record data](#) to manage use cases like getting metamodel on a custom programmatic enumeration, etc.

See also:

- `TableTrigger.handleNewContext`

See also:

- [Default value](#)

After being modified on the client side, the record can be submitted to be persisted using the built-in [Form insert operation](#) or [Insert operation](#) operations.

These prepare operations may use one of the following methods:

- `GET` HTTP method,
- `POST` HTTP method without message body or
- `POST` HTTP method with a message body and optionally [session parameters](#).

Available URL formats are:

- **Prepare for creation**

- **Record:**

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{tablePath}:prepareForCreation
```

- **Record field:**

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{tablePath}:prepareForCreation/{pathInRecord}
```

- **Prepare for duplication**

- **Record:**

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{tablePath}/
{encodedPrimaryKey}:prepareForDuplication
```

- **Record field:**

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{tablePath}/
{encodedPrimaryKey}:prepareForDuplication/{pathInRecord}
```

Where:

- `{category}` corresponds to the [operation category](#) (possible values are: `data` or `data-compact`).
- `{dataspace}` corresponds to `B` followed by the dataspace identifier.
- `{dataset}` corresponds to the dataset identifier.
- `{tablePath}` corresponds to the table path node.
- `{encodedPrimaryKey}` corresponds to the percent encoded representation of the primary key (see [RFC-3986 Uniform Resource Identifier](#)).
- `{pathInRecord}` corresponds to the path starting from the table node.

Parameters

The following parameters are applicable to the prepare operations:

PARAMETER	DESCRIPTION
<code>includeDetails</code>	Includes the <code>details</code> field in the metamodel and the content, for each indirect reachable resource. The returned value corresponds to its URL resource. Type <code>Boolean</code>, default value is <code>true</code>. See also: includeMetamodel
<code>includeMeta</code>	<i>Deprecated since 6.1.0, replaced by includeMetamodel.</i>
<code>includeMetamodel</code>	Includes the <code>meta</code> field holding the description of the structure returned in the <code>content</code> field. <code>Boolean</code> type, default value is <code>false</code>. See also: Metamodel
<code>includeSelector</code>	Includes the <code>selector</code> field in the content, for each indirect reachable resource. The returned value corresponds to its URL resource. Type <code>Boolean</code>, default value is <code>true</code>. See also: selector

Selector parameters

The available parameters are similar to the selector operation ones.

See also:

- [Selector parameters](#)

HTTP codes

HTTP CODE	DESCRIPTION
200 (<i>OK</i>)	The selected resource has been successfully retrieved.
400 (<i>Bad request</i>)	The request is incorrect. This occurs when: • The selected field in a record is sub-terminal, • <code>pageSize</code> parameter value is below 2, or is a string different from <code>unbounded</code>. • 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.
401 (<i>Unauthorized</i>)	Authentication has failed.
403 (<i>Forbidden</i>)	The selected resource is hidden from the authenticated user.
404 (<i>Not found</i>)	The selected resource could not be found.

Response body

After a successful prepare for a creation or duplication request, the transient record is returned in the response body. The content depends on the provided parameters and selected data. However, it takes a format similar to the select operation record.

Examples: [Extended JSON](#), [Compact JSON](#).

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/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}[:mass]
```

Where:

- **{category}** corresponds to the [operation category](#) (possible values are: **data** or **data-compact**).
- **{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.
- **:mass** extended action is required when updating a filtered set of records (see [filter parameter](#)).

Parameters

The following parameters are applicable with the insert operation.

PARAMETER	DESCRIPTION
includeDetails	Includes the details field in the content to access to the details of the data. The

PARAMETER	DESCRIPTION
	returned value corresponds to its URL resources. Type <code>Boolean</code>, the default value is <code>false</code>. Note: Only applicable on the record table. See also: • Insert operation report
<code>includeForeignKey</code>	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. <code>Boolean</code> type, the default value is <code>false</code>. Note: Only applicable on the record table. See also: • Insert operation report
<code>includeLabel</code>	Includes the <code>label</code> field in the answer for each record. Possible values are: • <code>yes</code>: the <code>label</code> field is included. • <code>no</code>: the <code>label</code> field is not included (use case: integration). <code>String</code> type, the default value is <code>no</code>. Note:

PARAMETER	DESCRIPTION
	Only applicable on the record table. See also: • Insert operation report
<code>updateOrInsert</code>	Specifies the behavior when the record to insert already exists: • If <code>true</code>: the existing record is updated with new data. For a request on a record table, the <code>code</code> field is added to the report in order to specify if this is an insert <code>201</code> or an update <code>204</code>. • If <code>false</code> (default value): a client error is returned and the operation is aborted. <code>Boolean</code> type value.
<code>byDelta</code>	Specifies the behavior for setting value of nodes that are not defined in the request body. This is described in the Update modes section. <code>Boolean</code> type, the default value is <code>true</code>. Note: Applicable on record in update mode and if the updateOrInsert parameter is <code>true</code>.
<code>filter</code>	Quick search predicate or complete XPath predicate expression that defines the records to which a massive update is applied. The request body must be a record structure with the fields to update. This parameter must imperatively be used with :mass

PARAMETER	DESCRIPTION
	extended action otherwise an error is returned. The ebx-all value is used to select all records. When the viewPublication query parameter is used, the scope is limited to the filtered view. <code>String</code> type value, it must not be empty otherwise an error would be returned. Note: When using the query parameters includeDetails, includeLabel or includeForeignKey, only the first 1000 report entries are returned in the report.
<code>viewPublication</code>	Specifies the name of the published view to be considered during a massive update. The behavior of this parameter is described in the section EBX® as a Web Component. <code>String</code> 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. <code>Boolean</code> type, default value is <code>false</code>. See Blocking and non-blocking constraints 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 [Extended JSON](#), [Compact JSON](#)).
- **Record table:** Similar to the select operation on a table but without the pagination information (example [Extended JSON](#), [Compact JSON](#)).

See also:

- [Inheritance](#)

HTTP codes

HTTP CODE	DESCRIPTION
200 (<i>OK</i>)	If the request relates to a record table. The insert request was applied successfully, an optional report is returned in the response body.
201 (<i>Created</i>)	If the request relates to a record. A new record has been created, in this case, the header field Location is returned with its resource URL.
204 (<i>No content</i>)	If the request relates to a record. Only available if updateOrInsert is true, an existing record has been successfully updated, in this case, the header field Location is returned with its resource URL.

HTTP CODE	DESCRIPTION
400 (<i>Bad request</i>)	The request is incorrect. This occurs when the body message structure does not comply with what was mentioned in Message body .
403 (<i>Forbidden</i>)	Authenticated user is not allowed to create a record 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, only available if <code>updateOrInsert</code> is <code>true</code> , the Optimistic locking is activated and the content has changed in the meantime, it must be reloaded before update.
422 (<i>Unprocessable entity</i>)	The request cannot be processed. This occurs when: - A blocking validation error occurs (only available if <code>blockingConstraintsDisabled</code> is <code>false</code>). - The record cannot be inserted because a record with the same primary key already exists (only available if <code>updateOrInsert</code> is <code>false</code>). - The record cannot be inserted because the definition of the primary key is either non-existent or incomplete. - The record cannot be updated because the value of the primary key cannot be modified.

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](#)). 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](#)

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/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}
```

- **Record:** modifies the content of selected record.

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}
```

Note:

Also available for `POST` HTTP method. In this case, the URL must point to the table and the parameter `updateOrInsert` must be set to `true`.

Note:

Use [Insert operation](#) with `updateOrInsert=true` and `byDelta=true` parameters to modify the content of several records.

- **Field:** update of a single field of the selected record.

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}/{pathInRecord}
```

Note:

The field must be either a terminal node or above.

Where:

- `{category}` corresponds to the [operation category](#) (possible values are: `data` or `data-compact`).
- `{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 percent encoded representation of the primary key (see [RFC-3986 Uniform Resource Identifier](#)).
- `{pathInRecord}` corresponds to the path starting from the table node.

Parameters

Here are the parameters applicable with the update operation.

PARAMETER	DESCRIPTION
<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. <code>Boolean</code> type, default value is <code>false</code> . See Blocking and non-blocking constraints for more information.
<code>byDelta</code>	Specifies the behavior for setting value of nodes that are not defined in the request body. This is described in the Update modes section. <code>Boolean</code> type, the default value is <code>true</code> .
<code>checkNotChangedSinceLastUpdateTime</code>	Timestamp in datetime format used to ensure that the record has not been modified since the last read. Also see the Optimistic locking section. <code>DateTime</code> type value.

Message body

The request must define a message body.

The structure is the same as for:

- the dataset node (sample [Extended JSON](#), [Compact JSON](#)),
- the record (sample [Extended JSON](#), [Compact JSON](#)),
- the record fields (sample [Extended JSON](#), [Compact JSON](#)),

depending on the updated scope, by only keeping the `content` entry.

See also:

- [Inheritance](#)

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 Optimistic locking 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: - A blocking validation error occurs (only available if <code>blockingConstraintsDisabled</code> is <code>false</code>). - The record cannot be updated because the value of the primary key cannot be modified.

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}[:mass]
```

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 percent encoded representation of the primary key (see [RFC-3986 Uniform Resource Identifier](#)).
- `:mass` extended action is required when deleting a filtered set of records (see [filter parameter](#)).

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](#)

Parameters

Here are the following parameters applicable with delete operation.

PARAMETER	DESCRIPTION
<code>includeOcculting</code>	Includes occulted records. <code>Boolean</code> type, the default value is <code>false</code> .
<code>inheritIfInOccultingMode</code>	<i>Deprecated since 5.8.1</i> While it remains available for backward compatibility reasons, it will eventually be removed in a future version. Inherits the record if it is in occulting mode. <code>Boolean</code> type, the default value is <code>true</code> .
<code>checkNotChangedSinceLastUpdateTime</code>	Timestamp in datetime format used to ensure that the record has not been modified since the last read. Also see the Optimistic locking section. <code>DateTime</code> 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. <code>Boolean</code> type, default value is <code>false</code> . See Blocking and non-blocking constraints for more information.
<code>filter</code>	Quick search predicate or complete XPath predicate expression that defines the records to which the massive delete is applied. This parameter must imperatively be used with <code>:mass</code> extended action otherwise an error would be returned. The ebx-all value is used to select all records. When the viewPublication query parameter is used, the scope is limited to the filtered view. <code>String</code> type value, it must not be empty otherwise an error would be returned.
<code>viewPublication</code>	Specifies the name of the published view to be considered during the massive delete. The behavior of this parameter is described in the

PARAMETER	DESCRIPTION
	section EBX® as a Web Component . <code>String</code> type value.

Message body

The request must define a message body only when deleting several records without using `filter` or `deleteAll` query parameter:

- **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`

Examples: [Extended JSON](#), [Compact JSON](#).

HTTP codes

HTTP CODE	DESCRIPTION
<code>200</code> (<i>OK</i>)	The operation has been executed successfully. A report is returned in the response body.
<code>400</code> (<i>Bad request</i>)	The request is incorrect. This occurs when: • the structure of the message body does not comply with Message body.

HTTP CODE	DESCRIPTION
	the message body contains a record table while the URL specifies a record.
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 Optimistic locking is activated and the content has changed in the meantime, it must be reloaded before deleting the record. The parameter value <code>checkNotChangedSinceLastUpdateTime</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**.

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/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}:count
```

Note:

Not applicable with the **history** category.

- **Table** depending on the **operation category**:

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
```

- **Field** depending on the **operation category**:

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
```

Note:

The field must be either an association node, a selection node, a terminal node or above.

Where:

- `{category}` corresponds to the [operation category](#) (possible values are: `data` or `data-compact`).
- `{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 percent encoded representation of the primary key (see [RFC-3986 Uniform Resource Identifier](#)).
- `{pathInRecord}` corresponds to the path starting from the table node.

Parameters

The following parameters are applicable to the count operation.

PARAMETER	DESCRIPTION
<code>count</code>	<i>Deprecated since 6.0.0</i> , replaced by the extended action from the URL. It is used to specify whether this is a count operation or a selection operation. <code>Boolean</code> type, default value is <code>false</code> .

Table parameters

The following parameters are applicable to tables, associations and selection nodes.

PARAMETER	DESCRIPTION
<code>filter</code>	Quick search predicate or complete XPath predicate expression defines the field values to which the request is applied. If empty, all records will be considered. <code>String</code> type value. Note:

PARAMETER	DESCRIPTION
	The history code operation value is usable with the <code>ebx-operationCode</code> path field from the <code>meta</code> section associated with this field.
<code>historyMode</code>	Specifies the filter context applied on table. <code>String</code> type, possible values are: <code>CurrentDataSpaceOnly</code> : history in current dataspace <code>CurrentDataSpaceAndAncestors</code> : history in current dataspace and ancestors <code>CurrentDataSpaceAndMergedChildren</code> : history in current dataspace and merged children <code>AllDataSpaces</code> : history in all dataspace The default value is <code>CurrentDataSpaceOnly</code>. See also: <code>history</code> Note: This parameter is ignored with the <code>data</code> category.
<code>includeOcculting</code>	Includes the records in occulting mode. <code>Boolean</code> type, default value is <code>false</code>.

PARAMETER	DESCRIPTION
<code>viewPublication</code>	Specifies the name of the published view to be considered during the count execution. This parameter can be combined with: the <code>filter</code> parameter as the logical <code>and</code> operation. The behavior of this parameter is described in the section EBX® as a Web Component. <code>String</code> type value.

Selector parameters

The following parameters are only applicable to fields that return an enumeration, foreign key or `osd:resource`.

PARAMETER	DESCRIPTION
<code>selector</code>	Specifies whether: <code>true</code>: returns the number of all possible values <code>false</code>: returns the number of possible values for the current field. <code>Boolean</code> type, default value is <code>false</code>. Note: This parameter is ignored with the <code>history</code> category.
<code>selectorFilter</code>	Specifies the filter of the selector. <code>String</code> type value, the syntax complies with the Quick Search.

HTTP codes

HTTP CODE	DESCRIPTION
200 (<i>OK</i>)	The selected resource is successfully counted.
400 (<i>Bad request</i>)	The request is incorrect. This occurs when: • The selected field in a record or a dataset is sub-terminal. • The selected dataset field is a dataset tree. • The quick search predicate or complete XPath predicate of the filter parameter is malformed or contains unfilterable nodes. • The table view for the viewPublication parameter is either hierarchical, non-existent or non-published. • The selector parameter is used for a non-enumerated node, or the firstElementIndex is negative, higher than or equal to the number of values.
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 `includeMetadata` parameter including the `system` value.

See [Technical data](#) for more information.

The `update_time` property value must be included in the following request. If the record has been changed since the specified time, the update or delete operation will be cancelled.

- **Record:** update whole or partial content of the selected record.

The `update_time` value should be added to the request body to prevent update of a modified record.

See the [JSON](#) example of a record.

- **Field:** update of a single field of the selected record.

The `update_time` value must be declared in the request URL by the `checkNotChangedSinceLastUpdateTime` parameter to prevent the update of a modified record.

The `update_time` property value can also be used in the request URL `checkNotChangedSinceLastUpdateTime` parameter to prevent deletion on a modified record.

Note:

The `checkNotChangedSinceLastUpdateTime` 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

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
<code>inheritance</code>	record or table metamodel	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. See also: <code>inheritance</code> property in metamodel.

PROPERTY	LOCATION	DESCRIPTION
<code>inheritedField</code>	field metamodel	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. See also: • <code>inheritedField</code> property in metamodel.
<code>inheritanceMode</code>	record in child dataset	Specifies the record's inheritance state. To set a record's inheritance from <code>overwrite</code> to <code>inherit</code>, its <code>inheritanceMode</code> value must be explicitly provided in the request. In this specific case, the <code>content</code> property will be ignored if present. <code>occult</code> and <code>root</code> explicit values are always ignored. An <code>overwrite</code> explicit value without a <code>content</code> property is ignored. Note: Inherited record's fields are necessarily <code>inherit</code>. Note:

PROPERTY	LOCATION	DESCRIPTION
		Root records in child dataset will always be <code>root</code>. Possible values are: <code>root</code>, <code>inherit</code>, <code>overwrite</code>, <code>occult</code>. For more information, see Record lookup mechanism.
	field in <code>overwrite</code> record	Specifies the field's inheritance state. To set a field's inheritance to <code>inherit</code>, its <code>inheritanceMode</code> value must be explicitly provided in the request. The <code>content</code> property will be ignored in this case. <code>overwrite</code> explicit value without a <code>content</code> property is ignored. Note: <code>inheritanceMode</code> at field level does not appear for <code>root</code>, <code>inherit</code> and <code>occult</code> records. Note: <code>inheritedFieldMode</code> and <code>inheritanceMode</code> properties cannot be both set on the same field. Possible values are: <code>inherit</code>, <code>overwrite</code>. For more information, see

PROPERTY	LOCATION	DESCRIPTION
		Inheritance and value resolution.
<code>inheritedFieldMode</code>	inherited field	Specifies the inherited field's inheritance state. To set a field's inheritance to <code>inherit</code>, its <code>inheritedFieldMode</code> value must be explicitly provided in the request. The <code>content</code> property will be ignored in this case. <code>overwrite</code> explicit values without a <code>content</code> property are ignored. Note: <code>inheritedFieldMode</code> and <code>inheritanceMode</code> properties cannot be both set on the same field. Note: <code>inheritedFieldMode</code> has priority over the <code>inheritanceMode</code> property. Possible values are: <code>inherit</code>, <code>overwrite</code>. For more information, see Value lookup mechanism.

Look up table views operation

The "look up published views" operation may use one of the following methods:

- `GET` HTTP method or

- **POST** HTTP method without message body.

The URL format is:

```
http[s]://<host>[:<port>]/<ebx-dataservices>/rest/data/v1/{dataspace}/
{dataset}/{tablePath: [&#94;:]*}:publishedViews
```

Where:

- **{dataspace}** corresponds to **B** followed by the dataspace identifier or to **V** followed by the image identifier.
- **{dataset}** corresponds to the dataset identifier.
- **{tablePath: [^:]*}** corresponds to the table path.

Parameters

No specific parameter for this operation.

Codes HTTP

HTTP CODE	DESCRIPTION
200 (<i>OK</i>)	Information successfully returned.
400 (<i>Bad request</i>)	Incorrect request, contains an error.
401 (<i>Unauthorized</i>)	Authentication has failed.
403 (<i>Forbidden</i>)	The specified resource read permission has been denied to the authenticated user.
404 (<i>Not found</i>)	The selected resource cannot be found.

Response body

It contains a collection of authorized view information (including the access URI).

Example: [JSON](#).

Form data operations

The form-data category operations concern the dataset fields, records or record fields when creating or modifying content must comply with constraints on user form submit. They are intended to be used in a user form management context.

The request body of an operation is very similar to the equivalent operation of the `data` category in the sense that an incoming data validation feature and a result report has been added to the response. The data validation can not be deactivated, a parameter

`blockingConstraintsDisabled` from the `data` category operations is not applicable. The validation phase fails when, at least, one constraint with a `blocksCommit` level set to `onInsertUpdateOrDelete` or `onUserSubmit-checkModifiedValues` is violated.

See [Blocking and non-blocking constraints](#) for more information.

Form insert operation

Insert form uses the `POST` HTTP method and requires a message body holding the data. This operation supports one or multiple records in a single transaction. Moreover, it is also possible to update record(s).

- **Record:** validate and insert a new record or modify an existing one in the selected table.
- **Record table:** validate and insert or modify one or more records in the selected table, while securing a consistent response. 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. The maximum number of records is limited to 100.

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/{dataspace}/{dataset}/{pathInDataset}[:mass]
```

Where:

- `{category}` corresponds to the [operation category](#) (possible values are: `form-data` or `form-data-compact`).
- `{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.
- `:mass` extended action is required when updating a filtered set of records (see [filter parameter](#)).

Parameters

The following parameters are applicable with this insert operation.

PARAMETER	DESCRIPTION
<code>includeDetails</code>	Includes the <code>details</code> field in the answer to access the data details. The returned value corresponds to its URL resources. Type <code>Boolean</code>, the default value is <code>false</code>. Note: Only applicable for a multiple records insertion.
<code>includeForeignKey</code>	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. <code>Boolean</code> type, the default value is <code>false</code>. Note: Only applicable for a multiple records insertion.
<code>includeLabel</code>	Includes the <code>label</code> field in the answer for each record. Possible values are: • <code>yes</code>: the <code>label</code> field is included.

PARAMETER	DESCRIPTION
	• <code>no</code> : the <code>label</code> field is not included (use case: integration). <code>String</code> type, the default value is <code>no</code>. Note: Only applicable for a multiple records insertion.
<code>updateOrInsert</code>	Specifies the behavior when the record to insert already exists: • If <code>true</code> : the existing record is updated with new data. • If <code>false</code> : a client error is returned and the operation is aborted. <code>Boolean</code> type, default value is <code>false</code>.
<code>filter</code>	Quick search predicate or complete XPath predicate expression that defines the records to which a massive update is applied. The request body must be a record structure with the fields to update. This parameter must imperatively be used with :mass extended action otherwise an error is returned. The ebx-all value is used to select all records. When the viewPublication query parameter is used, the scope is limited to the filtered view. <code>String</code> type value, it must not be empty otherwise an error would be returned.

PARAMETER	DESCRIPTION
	Note: When using the query parameters <code>includeDetails</code>, <code>includeLabel</code> or <code>includeForeignKey</code>, only the first 1000 report entries are returned in the report.
<code>viewPublication</code>	Specifies the name of the published view to be considered during a massive update. The behavior of this parameter is described in the section EBX® as a Web Component. <code>String</code> type value.

Message body

The request must define a message body. The format is similar to the `data` category insert operation's [message body](#).

HTTP codes

HTTP CODE	DESCRIPTION
<code>200</code> (<i>OK</i>)	- The request body holds multiple records: the insert/update request was applied successfully, a report is returned in the response body. - The request body holds only one record and <code>updateOrInsert</code> is <code>true</code>: the update request was applied successfully, a report is returned in the response body.
<code>201</code> (<i>Created</i>)	The request body holds only one non-existing record: a new record has been created and the header field <code>Location</code> is

HTTP CODE	DESCRIPTION
	returned with its resource URL. Moreover, a report is returned in the response body.
400 (<i>Bad request</i>)	The request is incorrect. This occurs when the body message structure: • does not comply with what was mentioned in Message body. • complies with what was mentioned in Message body, but number of record insert/update more than 100.
403 (<i>Forbidden</i>)	Authenticated user is not allowed to create a record 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, only available if <code>updateOrInsert</code> is <code>true</code>, the Optimistic locking is activated and the content has changed in the meantime, it must be reloaded before update.
422 (<i>Unprocessable entity</i>)	The request cannot be processed. This occurs when: • A blocking constraint has been violated. In this case an appropriate report is returned in the response body. • The record cannot be inserted because another one with the same primary key already exists (only available if <code>updateOrInsert</code> is <code>false</code>).

HTTP CODE	DESCRIPTION
	- The record cannot be inserted because the definition of the primary key is either non-existent or incomplete. - The record cannot be updated because the value of the primary key cannot be modified.

Response body

The response body will always hold a validation report. However in case of failure, the response body corresponds to a JSON `Exception handling` response.

- **Record:** The header field `Location` is returned with its URL resource. The validation report will be include in response body.
- **Record table:** (optional) Contains a list of validations report for each element, corresponding to the multiple validation report.

See also:

- [Form data category](#)

Form update operation

As for the `data` or `data-compact` category update operation, it allows the modification of a single dataset or record and the `PUT` HTTP method must be used. Available URL formats are:

- **Dataset node:** validates and modifies the values of terminal nodes contained in the selected node.

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}
```

- **Record:** validates and modifies the content of the selected record.

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}
```

Note:

Also available for `POST` HTTP method. In this case, the URL must point to the table and the parameter `updateOrInsert` must be set to `true`.

- **Field:** validates and update of a single field of the selected record.

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/{category}/v1/
{dataspace}/{dataset}/{pathInDataset}/{encodedPrimaryKey}/{pathInRecord}
```

Note:

The field must be either a terminal node or above.

Where:

- `{category}` corresponds to the [operation category](#) (possible values are: `form-data` or `form-data-compact`).
- `{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 percent encoded representation of the primary key (see [RFC-3986 Uniform Resource Identifier](#)).
- `{pathInRecord}` corresponds to the path starting from the table node.

Parameters

Here are the parameters applicable with the update operation.

PARAMETER	DESCRIPTION
<code>byDelta</code>	Specifies the behavior for setting value of nodes that are not defined in the request body. This is described in the Update modes section. <code>Boolean</code> type, the default value is <code>true</code> .
<code>checkNotChangedSinceLastUpdateTime</code>	Timestamp in datetime format used to ensure that the record has not been modified since the last read. Also see the Optimistic locking section. <code>DateTime</code> type value.

Message body

The request must define a message body. The format is the same as for the `data` category update operation [message body](#).

HTTP codes

HTTP CODE	DESCRIPTION
<code>200</code> (<i>Ok</i>)	The update request was applied successful and a report is returned in the response body.
<code>400</code> (<i>Bad request</i>)	The request is incorrect. This occurs when the body request structure does not comply.
<code>403</code> (<i>Forbidden</i>)	Authenticated user is not allowed to update the specified resource or the request body contains a read-only field.
<code>404</code> (<i>Not found</i>)	The selected resource is not found.
<code>409</code> (<i>Conflict</i>)	Concurrent modification, the Optimistic locking is activated and the content has

HTTP CODE	DESCRIPTION
	changed in the meantime, it must be reloaded before the update.
422 (<i>Unprocessable entity</i>)	The request cannot be processed. This occurs when: • A blocking constraint has been violated. In this case an appropriate report is returned in the response body. • The record cannot be updated because the value of the primary key cannot be modified.

Response body

The response body has the same format and behavior as the [Form Insert operation](#).

Beta feature: Dataspace operations

These operations perform a selection or life cycle management over dataspaces.

See also:

- [Dataspaces](#)

Beta feature: Select dataspaces or snapshots

Select operation may use one of the following methods: `GET` or `POST`. A pagination mechanism is always enabled.

URL formats are:

- **Root:** returns root dataspaces

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/data/v1/
```


- **Children:** returns children dataspace of a given dataspace

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/data/v1/
{dataspace}:children
```

- **Snapshots:** returns snapshots of a given dataspace

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/data/v1/
{dataspace}:snapshots
```

- **Information:** returns a dataspace, or a snapshot, information

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/data/v1/
{dataspace}:information
```

Where:

- `{dataspace}` corresponds to `B` followed by the dataspace's identifier or to `V` followed by the snapshot's identifier.

Parameters

The following query parameters are applicable to **Root**, **Children** and **Snapshots** operations.

PARAMETER	DESCRIPTION
<code>includeClosed</code>	Includes the closed dataspace to the selection. <code>Boolean</code> type, default value is <code>false</code> .
<code>includeAdministration</code>	Includes the administration dataspace to the selection. <code>Boolean</code> type, default value is <code>false</code> .

PARAMETER	DESCRIPTION
<code>pageRecordFilter</code>	Specifies a server side built filter, for pagination, pointing to a record. This filter is strongly linked to the <code>pageAction</code> value and should not be modified on the client side. The filter takes the form of a record XPath predicate expression used to figure out the pagination contexts. <code>String</code> type. See also: • Pagination
<code>pageAction</code>	Specifies the pagination action to perform from the identifier held by <code>pageRecordFilter</code>. <code>String</code> type, default value is <code>first</code>. The possible values are: • <code>first</code> • <code>previous</code> • <code>next</code> • <code>last</code> See also: • <code>RequestPagination</code>
<code>pageSize</code>	Specifies the maximum number of records per page. <code>Integer</code> type, default value is based on the user preferences. The value should be between <code>2</code> and <code>100</code>.

HTTP codes

HTTP CODE	DESCRIPTION
200 (<i>OK</i>)	The request has been successfully handled.
400 (<i>Bad request</i>)	The request is incorrect. This occurs when: pageAction parameter's value is not included in allowed values. pageRecordFilter is malformed or non-existent when selecting next or previous page. pageSize parameter's value is out of bound.
401 (<i>Unauthorized</i>)	Authentication has failed.
403 (<i>Forbidden</i>)	The selected resource is hidden from the authenticated user.
404 (<i>Not found</i>)	The selected resource could not be found.

Response body

After a successful selection, the result is returned in the response body. The content depends on the provided parameters and selected data.

The format is linked to the selected object type:

- For **Root**, **Children**, or **Snapshots**, see the [JSON](#) example.
- For **Information**, see the [JSON](#) example.

Beta feature: Create a child dataspace or a snapshot

Creates the dataspace or snapshot as specified. This operation use the **POST** method with a body request (with no specific query parameter).

See also:

- `Repository.createHome`

URL format are:

- **Dataspace:**

```
http[s]://<host>[:<port>]/.../data/v1/{dataspace}:createDataspace
```

- **Snapshot:**

```
http[s]://<host>[:<port>]/.../data/v1/{dataspace}:createSnapshot
```

Where:

- `{dataspace}` corresponds to `B` followed by the dataspace identifier.

Request body

The body specifies the features of the dataspace or the snapshot to create.

See also:

- `HomeCreationSpec`

See the [JSON](#) example.

HTTP codes

HTTP CODE	DESCRIPTION
<code>201</code> (<i>Created</i>)	A new dataspace or snapshot has been created, in this case, the header field <code>Location</code> is returned with its resource URL.
<code>400</code> (<i>Bad request</i>)	The request body is not well-formed or contains invalid content.

HTTP CODE	DESCRIPTION
401 (<i>Unauthorized</i>)	Authentication has failed.
403 (<i>Forbidden</i>)	Permission was denied to create the specified resource for the authenticated user.
404 (<i>Not found</i>)	The resource specified in the request body cannot be found.
500 (<i>Internal error</i>)	Unexpected error thrown by the application. Error details can usually be found in EBX® logs.

Beta feature: Locking a dataspace

Locks the specified dataspace. If the dataspace is already locked:

- by the same user, then the lock is kept,
- by another user and during the wait duration, then the lock is acquired,
- otherwise, the lock is rejected.

This operation uses the `POST` method and consumes `Content-Type` header set to:

- `application/x-www-form-urlencoded` : with HTTP parameters in the body or
- `application/json` : with HTTP parameters in the URL and [Session parameters](#) into JSON body.

When it succeeds, no response body is returned.

See also:

- [Permissions for locking or unlocking a dataspace](#)
- `LockSpec.lock`

The URL format is:

```
http[s]://<host>[:<port>]/.../data/v1/{dataspace}:lock
```

Where:

- `{dataspace}` corresponds to `B` followed by the dataspace identifier.

Parameters

The following query parameters are applicable to the operation.

PARAMETER	DESCRIPTION
<code>durationToWaitForLock</code>	When the dataspace is locked by another user, specifies the maximum duration to wait for acquiring the lock. The duration is specified in seconds. If the value is set to <code>0</code> , then the locking attempt is performed immediately. Due to several reasons, the wait duration can not exceed <code>60</code> seconds. Otherwise, the value is overwritten with its maximum value. <code>Integer</code> type, default value is <code>0</code> .

HTTP codes

HTTP CODE	DESCRIPTION
<code>204</code> (<i>No content</i>)	The request has been successfully handled but no response body is returned.
<code>400</code> (<i>Bad request</i>)	The request URL or body is not well-formed or contains invalid content.
<code>401</code> (<i>Unauthorized</i>)	Authentication has failed.
<code>403</code> (<i>Forbidden</i>)	Permission was denied to lock the specified resource for the authenticated user.
<code>404</code> (<i>Not found</i>)	The resource specified in the URL cannot be found.
<code>409</code> (<i>Conflict</i>)	The resource is already locked by another user, the lock is rejected after <code>durationToWaitForLock</code> parameter's value.

HTTP CODE	DESCRIPTION
500 (<i>Internal error</i>)	Unexpected error thrown by the application. Error details can usually be found in EBX® logs.

Beta feature: Unlocking a dataspace

Unlocks the specified dataspace. If the dataspace is:

- locked by the same user, then the lock is released,
- locked by another user, the current user is an administrator, and the `forceByAdministrator` query parameter is set to `true`, then the lock is released,
- not locked, the lock status is unchanged,
- otherwise, the unlock is rejected.

This operation uses the `POST` method and consumes `Content-Type` header set to:

- `application/x-www-form-urlencoded`: with HTTP parameters in the body or
- `application/json`: with HTTP parameters in the URL and [Session parameters](#) into JSON body.

When it succeeds, no response body is returned.

See also:

- [Permissions for locking or unlocking a dataspace](#)
- `LockSpec.unlock`

The URL format is:

```
http[s]://<host>[:<port>]/.../data/v1/{dataspace}:unlock
```

Where:

- `{dataspace}` corresponds to `B` followed by the dataspace identifier.

Parameters

The following query parameters are applicable to the operation.

PARAMETER	DESCRIPTION
<code>forceByAdministrator</code>	When the dataspace is locked by another user, an administrator can force the unlocking. <code>Boolean</code> type, default value is <code>false</code> .

HTTP codes

HTTP CODE	DESCRIPTION
<code>204</code> (<i>No content</i>)	The request has been successfully handled but no response body is returned.
<code>400</code> (<i>Bad request</i>)	The request URL or body is not well-formed or contains invalid content.
<code>401</code> (<i>Unauthorized</i>)	Authentication has failed.
<code>403</code> (<i>Forbidden</i>)	Permission was denied to unlock the specified resource for the authenticated user.
<code>404</code> (<i>Not found</i>)	The resource specified in the URL cannot be found.
<code>409</code> (<i>Conflict</i>)	The resource is already locked by another user, or the current user is administrator but <code>forceByAdministrator</code> parameter is <code>false</code> , the unlock is rejected.
<code>500</code> (<i>Internal error</i>)	Unexpected error thrown by the application. Error details can usually be found in EBX® logs.

Beta feature: Merge a dataspace

Merges the specified dataspace to its parent. It is possible to perform deletion, after the merge, on history and / or on data.

This operation uses the `POST` method and consumes `Content-Type` header set to:

- `application/x-www-form-urlencoded` : with HTTP parameters in the body or
- `application/json` : with HTTP parameters in the URL and [Session parameters](#) into JSON body.

When it succeeds, no response body is returned.

Note:

The merge decision step is bypassed for merges performed through data services. In such cases, the data in the child dataspace automatically overrides the data in its parent.

See also:

- `ProcedureContext.doMergeToParent`
- [Deleting data and history](#)

The URL format is:

```
http[s]://<host>[:<port>]/.../data/v1/{dataspace}:merge
```

Where:

- `{dataspace}` corresponds to `B` followed by the dataspace identifier.

Parameters

The following query parameters are applicable to the operation.

PARAMETER	DESCRIPTION
<code>deleteHistoryOnMerge</code>	Sets whether the history associated with the specified dataspace will be deleted upon merge. <code>Boolean</code> type, default value is <code>false</code> . When this parameter is not specified in the request, the default value is <code>false</code> . It

PARAMETER	DESCRIPTION
	is possible to redefine the default value by specifying the property <code>ebx.dataservices.historyDeletionOnCloseOrMerge.default</code> in the EBX® main configuration file.
<code>deleteDataOnMerge</code>	Sets whether the specified dataspace and its associated snapshots will be deleted upon merge. <code>Boolean</code> type, default value is <code>false</code>. When this parameter is not specified in the request, the default value is <code>false</code>. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.dataDeletionOnCloseOrMerge.default</code> in the EBX® main configuration file.

HTTP codes

HTTP CODE	DESCRIPTION
<code>204</code> (<i>No content</i>)	The request has been successfully handled but no response body is returned.
<code>400</code> (<i>Bad request</i>)	The request URL contains invalid content.
<code>401</code> (<i>Unauthorized</i>)	Authentication has failed.
<code>403</code> (<i>Forbidden</i>)	Permission was denied to merge the specified resource for the authenticated user.
<code>404</code> (<i>Not found</i>)	The resource specified in the URL cannot be found.
<code>500</code> (<i>Internal error</i>)	Unexpected error thrown by the application. Error details can usually be found in EBX® logs.

Beta feature: Close a dataspace or a snapshot

Closes the specified dataspace or snapshot. It is possible to perform deletion, after close, on history and / or on data.

This operation uses the `POST` method and consumes `Content-Type` header set to:

- `application/x-www-form-urlencoded`: with HTTP parameters in the body or
- `application/json`: with HTTP parameters in the URL and [Session parameters](#) into JSON body.

When it succeeds, no response body is returned.

See also:

- `Repository.closeHome`
- [Deleting data and history](#)

The URL format is:

```
http[s]://<host>[:<port>]/.../data/v1/{dataspace}:close
```

Where:

- `{dataspace}` corresponds to `B` followed by the dataspace identifier or to `V` followed by the snapshot identifier.

Parameters

The following query parameters are applicable to the operation.

PARAMETER	DESCRIPTION
<code>deleteHistoryOnClose</code>	Sets whether the history associated with the specified dataspace will be deleted upon close. <code>Boolean</code> type, default value is <code>false</code>. When this parameter is not specified in the request, the default value is <code>false</code>. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.historyDeletionOnCloseOrMerge.def</code> in the EBX® main configuration file.

PARAMETER	DESCRIPTION
<code>deleteDataOnClose</code>	Sets whether the specified dataspace and its associated snapshots will be deleted upon close. <code>Boolean</code> type, default value is <code>false</code>. When this parameter is not specified in the request, the default value is <code>false</code>. It is possible to redefine the default value by specifying the property <code>ebx.dataservices.dataDeletionOnCloseOrMerge.default</code> in the EBX® main configuration file.

HTTP codes

HTTP CODE	DESCRIPTION
<code>204</code> (<i>No content</i>)	The request has been successfully handled but no response body is returned.
<code>400</code> (<i>Bad request</i>)	The request URL or body is not well-formed or contains invalid content.
<code>401</code> (<i>Unauthorized</i>)	Authentication has failed.
<code>403</code> (<i>Forbidden</i>)	Permission was denied to close the specified resource for the authenticated user.
<code>404</code> (<i>Not found</i>)	The resource specified in the URL cannot be found.
<code>422</code> (<i>Unprocessable entity</i>)	A blocking constraint has been violated.
<code>500</code> (<i>Internal error</i>)	Unexpected error thrown by the application. Error details can usually be found in EBX® logs.

Health operations

Overview

Operations health helps monitor the EBX® server.

Started operation

Returns status code 200 (Ok) when EBX® is started and initialized, otherwise it returns 503 (Service unavailable).

The URL format is:

```
http[s]://<host>[:<port>]/.../health/v1/started
```

HTTP codes

HTTP CODE	DESCRIPTION
200 (OK)	The request was successfully handled and a response body is returned.
503 (Service unavailable)	EBX® is not yet started and initialized. A message is returned to help with diagnostics.

Response body

See the [JSON](#) example.

Check operation

Returns the status code 200 (Ok) when EBX® is started and checked, otherwise it returns 503 (Service unavailable).

The URL format is:

```
http[s]://<host>[:<port>]/.../health/v1/check
```

HTTP codes

HTTP CODE	DESCRIPTION
200 (OK)	The request was successfully handled and a response body is returned.
503 (Service unavailable)	EBX® is not yet started and initialized. A message is returned to help with diagnostics.

Response body

See the [JSON](#) example.

Stopped operation

Returns the status code 200 (Ok) when the EBX® repository is shutdown, otherwise it returns 503 (Service unavailable).

The URL format is:

```
http[s]://<host>[:<port>]/.../health/v1/stopped
```

HTTP codes

HTTP CODE	DESCRIPTION
200 (OK)	The request was successfully handled and a response body is returned.
503 (Service unavailable)	EBX® is not yet started and initialized. A message is returned to help with diagnostics.

Response body

See the [JSON](#) example.

OpenAPI operations

Overview

The `api` category operations comply with the OpenAPI specification 3.0.X to generate JSON documents. These documents facilitate development and consumption by structuring and describing the available REST built-in resources and the operations associated with them.

OpenAPI document HATEOAS links are available through the select data operations when using the [includeOpenApiDetails](#) query parameter.

The REST OpenAPI services permissions are globally defined in [Global permissions](#).

Making the proper types mapping for generating a client, in a specific language, is essential. The [OpenAPI format](#) can be used to validate the input or to map the value to a specific type, in the chosen programming language. See [Content of simple type](#) for more information.

Note:

Tools that do not support a specific format may default back to the `type` alone.

Generate OpenAPI document

The operation uses the `GET` or `POST` HTTP method to generate the OpenAPI document of a dataset table or a schema node.

The URL format is:

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/api/v1/{category}/v1/{dataspace}/{dataset}/{pathInDataset}
```

Where:

- `{category}` corresponds to an [operation category](#) among [data](#) or [form-data](#).
- `{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:
 - a table node,
 - a dataset terminal node or above.

Note:

The generated document will not depend on user permissions. The whole fields will be presented.

HTTP codes

HTTP CODE	DESCRIPTION
200 (OK)	The request has been successfully handled.
401 (Unauthorized)	Authentication has failed.
403 (Forbidden)	Permission was denied to read the specified resource.
404 (Not found)	The resource, specified in the URL, cannot be found.

Staging operations

The [Staging](#) REST API is designed to interact with source and target servers. It allows users to manage domains using export and import of staging archives.

The complete description of the services is available via OpenAPI on the following link:

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/api/v1/staging/v1
```

In the [development run mode](#), the [Swagger UI](#) is available on the following link.

```
http[s]://<host>[:<port>]/ebx-dataservices/rest/api/v1/staging/v1/ui
```

See also:

- [Visual documentation](#)

Domain management operations

The *domains* service allows the users to manage domains. The provided features are:

- Select the list of available domains.
- Select a domain by unique name.
- Create a new domain.
- Select the components of a domain.
- Add a component to a domain.

- Remove a component from a domain, the component is deleted.

Staging element navigation operations

The *elements* service is designed to display the whole elements of an instance in a hierarchical way. It is mainly designed for a UI interface that supports an hierarchical view with levels handling. The purpose of using query parameters instead of a structured body is to use links for easier navigation. The provided features are:

- Expands the content of an element.
- Expands the content of a specific type of elements under a selected element.

Archive manipulation operations

The *archives* service is designed for importing/exporting staging archives to/from a running instance in a single request. It is also used for uploading archives for asynchronous import. The provided features are:

- Export the archive of a domain.
- Import an archive to a running instance.
- Upload an archive for an asynchronous import.

The easiest way to export a staging archive via cURL

```
curl --request GET 'http[s]://<host>[:<port>]/ebx-dataservices/rest/staging/v1/archives/<domain>?login=<login>&password=<password>'
```

The easiest way to import a staging archive via cURL

```
curl --request POST 'http[s]://<host>[:<port>]/ebx-dataservices/rest/staging/v1/archives?login=<login>&password=<password>'

--header 'Content-Type: application/zip'

--data-binary '@/<archive-path-on-disk>'
```

Export operations


The *exports* service is designed to handle advanced export features. The archive is exported and stored on the server for a predefined period. The user can retrieve information about the archive, and it will be available to download it until it expires. The provided features are:

- Initiate an export.
- Get the status of an initiated export.
- Download an exported archive.

This endpoint can work in both synchronous and asynchronous modes.


Synchronous mode

A blocking mode that responds after exporting the whole archive on the server. The Archive will be immediately downloadable.



Asynchronous mode

A non-blocking mode that responds with an export ID and launches a task to export the archive. The client must poll the status endpoint to retrieve the status of the export. When the status changes to "Exported", the archive would be downloadable.



For consistency, any staging asynchronous job is aborted if it does not complete before 60 minutes of work.

Import operations

The *imports* service is designed to handle advanced import features. The archive is uploaded and stored on the server for a predefined period. The user can initiate an import asynchronously or retrieve information about an in-progress or completed import.


The archive is uploaded using the *archives* service with the option `asyncImport=true`.

The provided features are:

- Initiate an import by using an already uploaded archive.
- Get the status/report of an initiated import.

The archive name can be retrieved from an old import to initiate a re-import of the same archive with different options. It will be available until it expires.

This endpoint works only in asynchronous mode:



Limitations

General limitations

- Indexes, in the request URL `{pathInDataset}` or `{pathInRecord}`, are not supported.
- Dataset nodes and field operations applied to nodes that are sub-terminal are not supported.

See [Access properties](#) for more information about terminal nodes.

Compact format limitations

- [Inheritance](#), [tracking information](#), [session parameters](#) and [procedure context](#) are not handled in the [compact format](#). Use the [extended](#) one instead.
- `History category` is not supported in the [compact format](#).

Read operations

- Within the `selector`, the pagination context is limited to the `nextPage` property.
- When the `sortByRelevancy` parameter is activated, the ones below are ignored: [sort](#), [sortOnLabel](#), [sortPriority](#) and sort criteria defined through the [viewPublication](#).
- Within the `viewPublication` parameter, the hierarchical and tile views are 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](#) for more information.

- The prepare for creation operation does not support a dataset node.
- The [select dataspace operation](#) does not sort them by label.
- The [select snapshots operation](#) does not sort them by creation date.
- The [select snapshots operation](#) can not include the initial snapshots.

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](#) for more information.

Directory operations

- Changing or resetting a password for a user is not supported.

OpenAPI operations

- The document generation is not available through the user interface.
- The document generation REST services does not support the YAML format.
- Supports only [Data operations](#) and [Form data operations](#) except:
 - Select **Dataset tree**, **Dataset node** and **Field** operations.
 - Insert or delete multiple records in a single request.
 - Use of the HTTP header `content-type: x-www-form-urlencoded` to send query parameters in the body request.
- The '[Basic Authentication Scheme](#)' is the only described method.

JSON Formats

- [Introduction](#)
- [Extended](#)
- [Compact](#)
- [Common](#)
- [Others](#)

Introduction

Overview

The JSON (JavaScript Object Notation) is the data-interchange format used by TIBCO EBX®'s [RESTful operations](#).

This format is lightweight, self-describing and can be used to design UIs or to integrate EBX® in company's information system.

The following JSON formats are available:

- Using the [extended format](#), the data context is exhaustive and contains features to retrieve metamodel data, labels, validation report, inheritance data, URLs (OpenAPI, details, selector, historyDetails, etc.), history advanced data, except for association fields and selection nodes. However, these fields are reachable from the response through URL links called details included by default.
- Using the [compact format](#), the data context is limited to master data without metamodel data except links for enumerated fields.

See also:

- [Hiding a field in Data Services](#)

The amount of retrieved data is limited by a [pagination](#) mechanism which can be configured.

URL links allow reaching:

- Tables, records, dataset non-terminal nodes, foreign keys, resource fields through the `details` property. See [includeDetails](#) for more information.

- Possible values for foreign keys or enumerations, by activating the `selector` parameter. See `includeSelector` for more information.

See also:

- [Activation and configuration](#)

Note:

JSON data are always encoded with the UTF-8 charset.

Extended

- [Introduction](#)
- [Global structure](#)
 - [JSON Request body](#)
 - [JSON Response body](#)
- [Metamodel](#)
 - [Structure for table](#)
 - [Structure for field](#)
 - [Structure for foreign key field](#)
 - [Structure for the association field](#)
 - [Structure for the linked field](#)
 - [Structure for the enumeration dependencies](#)
- [Sort criteria](#)
- [Validation report](#)
- [Constraints](#)
 - [Constraint Types](#)
- [Content](#)

- Node value
- Content of simple type
- Content of group and list
- Selector

Introduction

The JSON extended format is used to retrieve master data, technical information and metadata. It is designed in an expanded way that allows to include several features such as validation, sorting and so on. To activate the extended format, the unsuffixed REST category, like `data` or `form-data`, must be used in the URL.

Global structure

JSON Request body

The request body is represented by a JSON `object` whose content varies according to the operation and the category.

Data category

The request body contains at least a `content` property which hold master data values.

- **Dataset node**

Specifies the target values of terminal nodes under the specified node. This request is used for the dataset node update operation.

```
{  
  "content": {  
    "nodeName1": {  
      "content": true  
    },  
    "nodeName2": {  
      "content": 2  
    }  
  }  
}
```



```

    },

    "nodeName3": {

        "content": "Hello"

    }

}

}

```

See also:

- [Update operation](#)

• Record

Specifies the target record content by setting the value of each field. For missing fields, the behavior depends on the request parameter `byDelta`. This structure is used for table record insert, record update or selector's select operation when local enumeration dependency field values are declared.

See also:

- [Inheritance](#)

Some technical data can be added beside the `content` property such as `lastUpdateDate`.

See also:

- [Optimistic locking](#)

```
{  
  
  ...  
  
  "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](#)
- [Update operation](#)

• 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](#)

```

{
  "content": [
    {
      "content": "CM"
    },
    {
      "content": "PI"
    }
  ]
}
```

```

    ]
}

```

See also:

- [Update operation](#)

• 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 `byDelta` parameter of the request. This structure is used upon insert or update 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](#)
- [Update operation](#)

• Record table to 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](#)

• Field

Specifies the target field content. This request is used for field update.

The request has the same structure as defined in [node value](#) by only keeping the `content` entry. Additional entries are simply ignored.

See also:

- [Update operation](#)

• 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](#)

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 the category.

Data category

The selection operation contains two different parts.

The first one named `meta` contains the exhaustive structure of the response.

The second, 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",

    "openApiDetails": "http://.../api/.../rootName",

    "content": {

        "settings": {

            "details": "http://.../rootName/settings",

            "openApiDetails": "http://.../api/.../rootName/settings",

            "content": {

                "weekTimeSheet": {

                    "details": "http://.../rootName/settings/settingA",

                    "openApiDetails": "http://.../api/.../rootName/settings/
settingA"

                },

                "vacationRequest": {

                    "details": "http://.../rootName/settings/settingB",

                    "openApiDetails": "http://.../api/.../rootName/settings/
settingB"

                }

            }

        },

        "table1": {

            "details": "http://.../rootName/table1",

            "openApiDetails": "http://.../api/.../rootName/table1"

        },

```

```

        "table2": {
            "details": "http://.../rootName/table2",
            "openApiDetails": "http://.../api/.../rootName/table2"
        }
    }
}
}
}
}

```

The `meta` and `validation` properties are optional.

See also:

- [Metamodel](#)
- [Validation report](#)
- [Select operation](#)

• Dataset node

Contains the list of terminal nodes under the specified node.

```

{
    "meta": {
        "fields": [
            {
                "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](#)

• Table

JSON `Object` containing the following properties:

- (Optional) The [table metamodel](#),
- (Optional) The sort criteria applied,
- (Optional) The table validation report,
- The `rows` property corresponding to a JSON `Array` of the selected records. Each record is represented by an JSON `Object`. If no record is selected then the JSON `Array` is empty.
- (Optional) the `pagination` property, containing [pagination](#) data.


```
{  
  
  "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  
  
        },  
  
        ...  
  
      }  
  
    },  
  
  ],  
  
}
```

```

...

{
  "label": "Alfred Dreyfus",
  "details": "http://.../root/individu/10",
  "content": {
    "id": {
      "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](#)

• Record

JSON `Object` containing:

- The label,
- (Optional) The record URL,
- (Optional) The [technical data](#) ,
- (Optional) The [table metamodel](#) ,
- (Optional) The record validation report,
- (Optional) The inheritance mode of the record, which can be: `root` , `inherit` , `overwrite` or `occult` .

See also:

- [Record lookup mechanism](#)
- [Inheritance](#)

- 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,

        "selector": "http://.../rootName/table1?selector=true",

        "validation": [

            {

                "level": "error",

                "message": "Field 'Parent' is mandatory."
            }
        ]
    }
}

```

```

        }

    ]

}

},

"validation": {

    ...

}

}

```

See also:

- [Select operation](#)
- [Prepare operations](#)

• Fields

For association or selection nodes, contains the target table with the associated records if, and only if, the `includeDetails` parameter is set to `true`.

For other kinds of nodes, contains the current [node value](#).

See also:

- [Select operation](#)

• 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](#)

Note:

Nodes, records and fields, property and values may be hidden depending on their resolved permissions (see [permissions](#)) .

Metamodel

This section can be activated on demand with the `includeMetamodel` 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](#)
- [Prepare operations](#)

Structure for table

Table metamodel is represented by a JSON `Object` with the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>name</code>	<code>String</code>	Specifies the name of the authorized table defined in the model.	Yes
<code>label</code>	<code>String</code>	Specifies the table's label. If undefined, the name of the schema node is returned.	Yes
<code>description</code>	<code>String</code>	Specifies the table's description.	No
<code>type</code>	<code>String</code>	Specifies the node's type, the value is always equal to: <code>table</code> .	Yes
<code>minOccurs</code>	<code>Number</code>	Specifies the number of minimum authorized record(s).	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>maxOccurs</code>	<code>Number</code> or <code>String</code>	Specifies the number of maximum authorized record(s) or <code>unbounded</code> .	Yes
<code>history</code>	<code>Boolean</code>	Specifies if the table content is historized. Its value is <code>true</code> if history is activated, <code>false</code> otherwise. See also: • History	No
<code>primaryKeyFields</code>	<code>Array</code>	Specifies the primary key composition which is represented as an array of the paths.	Yes
<code>inheritance</code>	<code>Boolean</code>	Specifies whether the dataset inheritance is activated for the table. Its value is <code>true</code> if inheritance is activated, <code>false</code> otherwise. See also:	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		Inheritance and value resolution	
fields	Array	Specifies the direct children of the record node which are represented as an array of fields. Each field may also recursively contain sub-fields.	Yes

Structure for field

Each authorized field is represented by a JSON `Object` with the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
name	String	Specifies the name of the authorized field defined in the model.	Yes
label	String	Specifies the node's label. If undefined, the name of the schema node is returned.	Yes
description	String	Specifies the field's description.	No
type	String	Specifies the field's type, which can be a: <code>simple type</code> , <code>group</code> , <code>table</code> ,	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		foreignKey , association , etc.	
<code>minOccurs</code>	<code>Number</code>	Specifies the number of minimum authorized occurrence(s).	Yes
<code>maxOccurs</code>	<code>Number</code> or <code>String</code>	Specifies the number of maximum authorized occurrence(s) or <code>unbounded</code> .	Yes
<code>readOnly</code>	<code>Boolean</code>	Specifies whether the field access permission is defined as read-only from the data model access properties. That is, if its data can be read or written. Its value is <code>true</code> if the field is read only, <code>false</code> otherwise. See also: • Access properties • <code>AccessPermission</code>	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>autoIncrement</code>	<code>Boolean</code>	Specifies whether the field is auto-incremented. This property is only available for integer field type. Its value is <code>true</code> if the field is auto-incremented, <code>false</code> otherwise. See also: • Auto-incremented values	No
<code>checkNullInput</code>	<code>Boolean</code>	Specifies whether the check of the constraints on <code>null</code> fields, at user input, is activated or not. Its value is <code>true</code> when activated, <code>false</code> otherwise. The constraint check is activated when the field has the property <code>osd:checkNullInput</code> set to <code>true</code>.	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>inheritedField</code>	<code>Object</code>	Holds information related to the inherited field's value source. <pre>{ "inheritedField": { "sourceRecord": "/path/to/record", // (optional) "sourceNode": "./path/to/Node" } }</pre> See also: Inheritance and value resolution	No
<code>foreignKey</code>	<code>Object</code>	Holds information related to the target table. See also: Structure for foreign key field	No (*)
<code>association</code>	<code>Object</code>	Holds information related to the association table. See also: Structure for the association field	No (*)

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>enumeration</code>	<code>String</code>	Specifies if the field is an enumeration value. Possible values are: <code>dynamic</code> <code>foreignKey</code> <code>nomenclature</code> <code>programmatic</code> <code>resource</code> <code>static</code> See also: <code>SchemaFacetEnumeration</code> According to its value, indicates if the field can be used for retrieving possible values by using the <code>selector</code> request parameter.	No
<code>enumerationDependencies</code>	<code>Array</code> of <code>JSON Object</code>	Specifies the metadata enumeration dependencies list. Only available for enumeration types among: <code>foreignKey</code>, <code>dynamic</code> or <code>programmatic</code>.	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		See also: • Structure for the enumeration dependencies	
<code>valueFunction</code>	<code>Boolean</code>	Specifies if the field is a computed value. See also: • Computed values	No
<code>linkedField</code>	<code>Object</code>	Holds information related to the linked field. See also: • SchemaNode.getSchemaLinkedField() • Structure for the linked field	No
<code>pathInDataset</code>	<code>String</code>	Specifies the relative field's path starting from the schema node.	No (**)
<code>pathInRecord</code>	<code>String</code>	Specifies the relative field's path starting from the table node.	No (*)
<code>filterable</code>	<code>Boolean</code>	Specifies whether the field can be	No (*)

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		used to filter records using the <code>filter</code> request parameter.	
<code>hiddenFilterPolicy</code>	<code>String</code>	Specifies the hidden filter policy. The possible values are: <code>textSearchOnly</code>, <code>true</code> or <code>false</code>. See also: Hiding a field in structured search tools <code>SchemaNodeDefaultView.getHiddenInSearch</code>	No (*)
<code>sortable</code>	<code>Boolean</code>	Specifies whether the field can be used in sort criteria using the <code>sort</code> request parameter.	No (*)
<code>constraints</code>	<code>Array</code> of <code>JSON Object</code>	Specifies the field's non-programmatic constraints. See also: Constraints	No
<code>fields</code>	<code>Array</code> of <code>JSON Object</code>	Holds the structure and typing of each field group.	No

(*) Only available for table, record and record field operations.

(**) Only available for dataset tree operations.

Structure for foreign key field

The foreign key field metamodel is represented by a JSON `Object`.

```
{
  "dataspace": "BAuthors",
  "dataset": "Authors",
  "tablePath": "/root/Authors",
  "details": "http://.../BAuthors/Authors/root/Authors"
}
```

It holds the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>dataspace</code>	<code>String</code>	Specifies the target dataspace or snapshot identifier. If not specifies then it corresponds to the same dataspace as the table's one.	No
<code>dataset</code>	<code>String</code>	Specifies the target dataset's identifier. If not specifies then it corresponds to the same dataset as the table's one.	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>tablePath</code>	<code>String</code>	Specifies the target table's path.	Yes
<code>details</code>	<code>String</code>	Specifies the target table's REST resource URL. See also: <code>includeDetails</code> parameter	No
<code>historyDetails</code>	<code>String</code>	Specifies the history target table's REST resource URL. See also: <code>includeDetails</code> parameter	No

Structure for the association field

The association field metamodel is represented by a JSON `Object`. This object holds properties which depend on the association type:

- `tableRefInverse`

```
{
  "type": "tableRefInverse",
  "dataspace": "BTitles",
  "dataset": "Titles",
  "tablePath": "/root/Titles",
  "details": "http://.../BTitles/Titles/root/Titles",
```

```

    "fieldToSource": "/root/Titles/au_id",

    "filter": "./unit_price > 10"

}

```

- `linkTable`

```

{

    "type": "linkTable",

    "tablePath": "/root/Inventory",

    "details": "http://.../BInventory/Inventory/root/Inventory",

    "linkTablePath": "/root/Inventory",

    "fieldToSource": "/root/Inventory/store",

    "fieldToTarget": "/root/Inventory/item",

    "filter": "./price > 10"

}

```

- `xpathLink`

```

{

    "type": "xpathLink",

    "tablePath": "/root/Inventory",

    "details": "http://.../BAuthors/Authors/root/Inventory",

    "predicate": "/root/Inventory[./price < 100]",

    "filter": "./price > 10"

}

```

}

It holds the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>type</code>	<code>String</code>	Specifies the association type. The possible values are: <code>tableRefInverse</code>, <code>linkTable</code>, <code>xpathLink</code> See also: Associations	Yes
<code>dataspace</code>	<code>String</code>	Specifies the target dataspace or snapshot identifier. If not specifies, then it corresponds to the same dataspace as the table's one.	No
<code>dataset</code>	<code>String</code>	Specifies the target dataset's identifier. If not specifies, then it corresponds to the same dataset as the table's one.	No
<code>tablePath</code>	<code>String</code>	Specifies the path of the target table.	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>details</code>	<code>String</code>	Specifies the target table's REST resource URL. See also: • <code>includeDetails</code> parameter	No
<code>fieldToSource</code>	<code>String</code>	Specifies the field which refers to the source table of the association. Defined if the association type is <code>tableRefInverse</code> or <code>linkTable</code>.	No
<code>fieldToTarget</code>	<code>String</code>	Specifies the path of the field to the target table. Defined if the association type is <code>linkTable</code>.	No
<code>linkTablePath</code>	<code>String</code>	Specifies the path of the link table. Defined if the association type is <code>linkTable</code>.	No
<code>predicate</code>	<code>String</code>	Specifies the criteria of the association,	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		relative to the current node. Defined if the association type is <code>xpathLink</code> .	
<code>filter</code>	<code>String</code>	Specifies the XPath predicate expression to filter the associated objects.	No

Structure for the linked field

The linked field metamodel is represented by a JSON `Object`.

```
{
  "relationshipField": "/au_id",
  "tablePath": "/root/Authors",
  "linkedFieldPath": "/birth_date",
  "details": "http://.../BAuthors/Authors/root/Authors"
}
```

It holds the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>relationshipField</code>	<code>String</code>	Specifies the field's path holding the relationship used by this linked field.	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		See also: <code>SchemaLinkedField.getRelationshipNode</code>	
<code>tablePath</code>	<code>String</code>	Specifies the table's path referred by this linked field.	Yes
<code>linkedFieldPath</code>	<code>String</code>	Specifies the path referred by this linked field. See also: <code>SchemaLinkedField.getLinkedField</code>	Yes
<code>details</code>	<code>URI</code>	Specifies the target table's REST resource URL. See also: <code>includeDetails</code> parameter	No

See also:

- [Linked fields](#)
- `SchemaLinkedField`

Structure for the enumeration dependencies

The metamodel enumeration dependency object is represented by a JSON `Object`.

```
{
  "localModify": false,
  "dataspace": "BAuthors",
  "dataset": "Authors",
  "targetPath": "/root/Authors",
  "details": "http://.../BAuthors/Authors/root/Authors"
}
```

It holds the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>localModify</code>	<code>Boolean</code>	Specifies whether the enumeration dependency field is impacted by modifications performed on a field in the same record.	Yes
<code>dataspace</code>	<code>String</code>	Specifies the target dataspace or snapshot identifier of the enumeration dependency. If not specified then it corresponds to the same dataspace as the record's one.	No
<code>dataset</code>	<code>String</code>	Specifies the target dataset's identifier of the enumeration	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		dependency. If not specifies then it corresponds to the same dataset as the record's one.	
<code>targetPath</code>	<code>String</code>	Specifies the target path of the enumeration dependency.	Yes
<code>details</code>	<code>String</code>	Specifies the target REST resource table URL of the enumeration dependency. See also: • <code>includeDetails</code> parameter	No

Sort criteria

The sort criteria, applied to the request, can be returned on demand, by using the

`includeSortCriteria` parameter (deactivated by default). If it is activated, a `sortCriteria` property is directly added to the response root node.

A `sortCriteria` property is represented as a JSON `Array` that contains ordered sort criterion. Each sort criterion corresponds to a JSON `Object` with the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>path</code>	<code>String</code>	Field's path.	Yes
<code>order</code>	<code>String</code>	Possible values are: <code>asc</code> , <code>lasc</code> , <code>desc</code> or <code>ldesc</code> .	Yes

The `sortByRelevancy` property is represented by a JSON `Array`, itself holding a JSON `Object` with the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>order</code>	<code>String</code>	Defines whether the sort is done in ascending or descending order. Possible values are: <code>1asc</code> or <code>1desc</code> . See sortByRelevancy for more information.	Yes

Validation report

The validation can be activated on demand with the `includeValidation` parameter (deactivated by default). If it is activated, `validation` properties are directly added on target nodes with one or several messages. For messages without target node's path, a `validation` property is added on the root node.

A `validation` property is represented by a JSON `Array`, holding a JSON `Object` per message, corresponding to a validation item, with the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>level</code>	<code>String</code>	Severity level of the validation item. The possible values are: <code>info</code> , <code>warn</code> , <code>error</code> , <code>fatal</code> .	Yes
<code>message</code>	<code>String</code>	Description of the validation item.	Yes
<code>details</code>	<code>String</code> corresponding to an absolute URL.	URL of the resource associated with the validation item.	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		Only available on table and dataset scopes, if the associated resource exist and if it is included. See also: <code>includeDetails</code> parameter	

Constraints

This section is automatically activated when the node has a read-write permission. It provides the non-programmatic constraints declared on the data model.

The `constraints` property is represented by a JSON `Array`. Every constraint is taking the form of a JSON `Object`. They have the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>type</code>	<code>String</code>	Specifies the type of the constraint See also: Constraint Types	Yes
<code>level</code>	<code>String</code>	Specifies the severity level of the constraint. The possible values are: <code>fatal</code>, <code>error</code>, <code>warning</code> or <code>info</code>.	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>content</code>	<code>Object</code>	Specifies the value of the non-programmatic constraint. This value can be a simple type, like a <code>String</code> , or a JSON <code>Object</code> .	No
<code>min</code>	<code>Integer</code>	Specifies the minimum value for <code>excludeSegment</code> constraint type.	No
<code>max</code>	<code>Integer</code>	Specifies the maximum value for <code>excludeSegment</code> constraint type.	No
<code>message</code>	<code>String</code>	Specifies the validation message returned when the value of the field does not comply with the constraint.	Yes
<code>blocksCommit</code>	<code>String</code>	Specifies the control policy management for blocking or non-blocking constraint. See Blocking and non-blocking constraints for more information.	No

Constraint Types

This section lists the non-programmatic constraints types.

TYPE	DESCRIPTION
<code>mandatory</code>	Defines the field is mandatory
<code>fixedLengthStatic</code>	Defines the exact number of characters required for this field.
<code>fixedLengthDynamic</code>	Points out a field which provides the exact number of characters required for this one.
<code>minLengthStatic</code>	Defines the minimum number of characters allowed for this field.
<code>minLengthDynamic</code>	Points out a field which provides the minimum number of characters allowed for this one.
<code>maxLengthStatic</code>	Defines the maximum number of characters allowed for this field.
<code>maxLengthDynamic</code>	Points out a field which provides the maximum number of characters allowed for this field.
<code>excludeSegment</code>	Defines an inclusive range of not allowed values for this field.
<code>excludeValue</code>	Defines a list of not allowed values for this field.
<code>fractionDigits</code>	Defines the maximum number of decimal allowed for this field.
<code>pattern</code>	Defines the regular expression pattern that must be match by the field's value.
<code>totalDigits</code>	Defines the maximum number of digits allowed for this integer or decimal field.
<code>boundaryMinInclusiveStatic</code>	Defines the minimum value allowed (including the minimum value) for this field.

TYPE	DESCRIPTION
<code>boundaryMinExclusiveStatic</code>	Defines the minimum value allowed (excluding the minimum value) for this field.
<code>boundaryMinInclusiveDynamic</code>	Defines a field that provide the minimum value allowed (including the minimum value) for this field.
<code>boundaryMinExclusiveDynamic</code>	Defines a field that provide the minimum value allowed (excluding the minimum value) for this field.
<code>boundaryMaxInclusiveStatic</code>	Defines the maximum value allowed (including maximum value) for this field.
<code>boundaryMaxExclusiveStatic</code>	Defines the maximum value allowed (excluding the maximum value) for this field.
<code>boundaryMaxInclusiveDynamic</code>	Defines a field that provide the maximum value allowed (including the maximum value) for this field.
<code>boundaryMaxExclusiveDynamic</code>	Defines a field that provide the maximum value allowed (excluding the maximum value) for this field.

See also:

- [XML Schema supported facets](#)
- [Extended facets](#)

Content

This section can be deactivated on demand with the `includeContent` 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
<code>content</code>	Content of simple type Content of group and list	Contains the node value. Available for all nodes except for <code>association</code> and <code>selection</code>. However, their content can be retrieved by invoking the URL provided in their <code>details</code> property.	No
<code>details</code>	<code>String</code> corresponding to an absolute REST resource URL.	Returns the node details when invoked. The response type depends on the meta type. <code>foreignKey</code>: target record (available on table, record and field operations). <code>resource</code>: target resource (available on dataset node, table, record and field operations).	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		• <code>association</code> : target table containing the associated records (available on table and record operations). • <code>selection</code> : target table containing the associated records (available on table and record operations). • <code>group</code> : target dataset group node (available on dataset tree operation). • <code>table</code> : target table (available on dataset tree operation). Example: <pre>http://.../ BReference/ dataset/root/ table/pk/ associationField</pre>	

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>historyDetails</code>	<code>String</code> corresponding to an absolute REST resource URL.	Returns the node history details when invoked. See also: • <code>includeHistory</code>	<code>No</code>
<code>label</code>	<code>String</code>	Contains the foreign key or enumeration label in the current locale. The default label is returned if the current locale is not supported.	<code>No</code>
<code>inheritanceMode</code>	<code>String</code>	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. See also: • <code>inheritedFieldMode</code> • Record lookup mechanism	<code>No</code>

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>inheritedFieldMode</code>	<code>String</code>	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. See also: <code>inheritanceMode</code> Value lookup mechanism	<code>No</code>
<code>selector</code>	<code>String</code> corresponding to an absolute URL.	Contains the appropriate URL for selector operation. Example: <pre>http://.../ BReference/ dataset/root/ table/pk/ enumField? selector=true</pre>	<code>No</code>

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>validation</code>	<code>Array</code>	Contains the validation report that concerns the current node context. See also: Validation report	<code>No</code>

Content of simple type

A simple field value is stored in a JSON `Object` and its content is the value of the `content` property.

See also:

- [Simple types formats](#)

Content of group and list

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE
Group <code>xs:complexType</code>	<code>Object</code> Contains a property per sub-node.	Example for a simple-occurrence group. <pre>{ "road": { "content": "11 rue scribe" }, "zipcode": {</pre>	<code>group</code>

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE
		<pre> "content": "75009" }, "country": { "content": "France" } } </pre>	
List <code>maxOccurs > 1</code>	Array Contains an array of all field occurrences represented by a JSON Object. Each object is represented as a node value.	Example for a multi-occurrence field of <code>xs:int</code> type. <pre> [{ "content": 0 }, { "content": 1 }, { "content": 2 }, { "content": 3 }] </pre>	Meta of simple type, or group

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE
		] Example for a multi-occurrence group. [{ "content": { "road": { "content": "11 rue scribe" }, "zipcode": { "content": "75009" }, "country": { "content": "France" } } }, { "content": { "road": { "content": "711 Atlantic Ave" },	

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE
		<pre> "zipcode": { "content": "MA 02111" }, "country": { "content": "United States" } } }]</pre>	

Selector

By invoking the URL represented by the `selector` property on a field that provides an enumeration, it returns a JSON `Object` containing the following properties:

- `rows` corresponding to an `Array` of JSON `Object` where each one contains two entries: the returned `content` that can be persisted and the corresponding `label`. The list of established possible items depends on the current context.
- (Optional) `pagination` containing [pagination](#) data (activated by default).

```

{
  "rows": [
    {
      "content": "F",
      "label": "feminine"
    },

```



```
{
  "content": "M",
  "label": "masculine"
},
"pagination": {
  "nextPage": null
}
```

See also:

- [includeSelector](#)

Compact

- [Introduction](#)
- [Global structure](#)
 - [JSON Request body](#)
 - [JSON Response body](#)
- [Content](#)
 - [Node value](#)
 - [Content of simple types](#)
 - [Content of group, list and enumeration](#)

Introduction

The JSON compact format purpose is to retrieve the master data using a lightweight structure. It follows the key-value design to display data in the simplest way and without any technical information. To activate the compact format, the `compact` suffixed REST category, like `data-compact` or `form-data-compact`, must be used in URL.

Global structure

JSON Request body

The request body is represented by a JSON `Object` whose content varies according to the operation and the category.

Data category

- **Dataset node**

Specifies the target values of terminal nodes under the specified node. This request is used for the dataset node update operation.

```
{  
  "nodeName1": true,  
  "nodeName2": 2,  
  "nodeName3": "Hello"  
}
```

To ensure a JSON symmetric structure between the HTTP request and the HTTP response, enumeration node request format is like the following:

```
{  
  "enumerationNode": {  
    "key": "a key value"  
  }  
}
```

```
}
```

See also:

- [Update operation](#)

• Record

Specifies the target record content by setting the value of each field. For missing fields, the behavior depends on the request parameter `byDelta`. This structure is used for table record insert, record update or selector's select operation when local enumeration dependency field values are declared.

```
{
  "gender": {
    "key": "M"
  },
  "lastName": "Chopin",
  "lastName-en": "Chopin",
  "firstName": "Fryderyk",
  "firstName-en": "Frédéric",
  "birthDate": "1810-03-01",
  "deathDate": "1849-10-17",
  "jobs": [
    {
      "key": "CM"
```

```

    },

    {
        "key": "PI"
    }

],

"infos": [

    "https://en.wikipedia.org/wiki/Chopin"

]

}

```

See also:

- [Insert operation](#)
- [Update operation](#)

• 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.

```

{

    "jobs": [

        {

            "key": "CM"

        },

    ],

}

```

```

    {
      "key": "PI"
    }
  ]
}

```

See also:

- [Update operation](#)

• 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 records in the table.

```

[
  {
    "gender": {
      "key": "M"
    },
    "lastName": "Saint-Saëns",
    "firstName": "Camille",
    "birthDate": "1835-10-09",
    ...
  },

```

```

{
  "gender": {
    "key": "M"
  },
  "lastName": "Debussy",
  "firstName": "Claude",
  "birthDate": "1862-10-22",
  ...
}
]

```

See also:

- [Insert operation](#)
- [Update operation](#)

• Record table to delete

Defines one or more records. This structure is used upon deleting several records from the same table.

```

[
  {
    "details": "http://.../root/table/1"
  },
  {

```

```

    "details": "http://.../root/table/2"

  },

  {

    "primaryKey": "./oid=3"

  },

  {

    "foreignKey": "4"

  },

  ...

]

```

See also:

- [Delete operation](#)

• Field

Specifies the target field content. This request is used upon field update.

The request has the same structure as defined in [node value](#).

See also:

- [Update operation](#)

• Open or close user interface

Specifies whether the user interface is open or close and the unavailability message.

```
{
  "toolStatus": {
    "key": true,      // or false
    "label": "Open"
  },
  "toolStatusCloseMessage": "Access is temporarily forbidden for
maintenance."
}
```

See also:

- [User interface operations](#)

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 the authenticated user.

JSON Response body

The response body is represented by a JSON `Object` whose content depends on the operation and the category.

Data category

- **Dataset node**

Contains the list of terminal nodes under the specified node.


```
{
  "nodeName1": true,
  "nodeName2": 2,
  "nodeName3": "Hello"
}
```

See also:

- [Select operation](#)

• Table

JSON `Object` containing the following properties:

- `rows` corresponding to JSON `Array` of selected records. Each record is represented by a JSON `Object`. If no record is selected, then the JSON `Array` is empty.
- (Optional) `pagination` containing [pagination](#) data.

```
{
  "rows": [
    {
      "id": 1,
      "firstName": "Claude",
      "lastName": "Levi-Strauss"
    },
    {
      "id": 2,
```

```

        "firstName": "Sigmoud",

        "lastName": "Freud"

    },

    {

        "id": 3,

        "firstName": "Alfred",

        "lastName": "Dreyfus"

    }

],

"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](#)

• Record

JSON `Object` containing the record content.

```
{
  "pk": "pk1",
  "name": "Name1",
  "name-fr": "Name1",
  "parent": null
}
```

See also:

- [Select operation](#)
- [Prepare operations](#)

• Fields

The compact json format does not support the association and selection nodes.

For other kinds of nodes, they contain the current [node value](#).

See also:

- [Select operation](#)

• Retrieve the user interface state

Contains the user interface status and the unavailability message.

```
{
  "toolStatus": {
```

```

    "key": true,      // or false

    "label": "Open"

  },

  "toolStatusCloseMessage": "Access is temporarily forbidden for
maintenance."

}

```

See also:

- [User interface operations](#)

Note:

Nodes, records and fields property and value may be hidden depending on their resolved permissions (see [permissions](#)) .

Content

This section is always included and contains master data without any additional fields.

Node value

The node value contains only the data or the label and the details link in case of enumeration. It is available for all nodes except association and selection.

Content of simple types

Corresponds to a key-value JSON entry which describes the node content.

See also:

- [Simple types formats](#)

Content of group, list and enumeration

XML SCHEMA	JSON FORMAT	EXAMPLES
Group <code>xs:complexType</code>	<code>Object</code> Contains a property per sub-node.	Example for a simple-occurrence group. <pre>{ "road": "11 rue scribe", "zipcode": "75009", "country": "France" }</pre>
List <code>maxOccurs > 1</code>	<code>Array</code> Contains an array of all field occurrences represented by a JSON <code>Object</code> or a simple type. Each JSON <code>Object</code> is composed of node values .	Example for a multi-occurrence field of <code>xs:int</code> type. <pre>[0, 1, 2, 3, 4]</pre> Example for a multi-occurrence group. <pre>[{ "road": "11 rue scribe",</pre>

XML SCHEMA	JSON FORMAT	EXAMPLES
		<pre>"zipcode": "75009", "country": "France" }, { "road": "711 Atlantic Ave", "zipcode": "MA 02111", "country": "United States" }]</pre>
Enumeration <code>xs:string</code>	<code>Object</code> Contains <code>key</code> , <code>link</code> and <code>label</code> properties.	Example for a foreign key. <pre>{ "key": "1", "details": "http://.../Bdataspace/ "label": "Française" }</pre>

Common

- [Introduction](#)
- [Global structure](#)
 - [JSON Response body](#)
- [Form validation report](#)
- [Content](#)

- Pagination
- Content of simple type
- Insert operation report
- Delete operation report
- Technical data
- Update modes
- Known limitations
 - Field values
 - Indexing and search strategy

Introduction

The common specifications are used in both [compact](#) and [Extended](#) formats.

Global structure

JSON Response body

Form data category

The response body contains a JSON `Object` per handled resources. That is, the JSON root `Object` for a single handled resource is a JSON `Object` as well as for multiple handled resources. The several properties are directly placed into the JSON `Object`.

- If the insert/update of a single record request or update field in table/dataset request was successful, the `code` JSON property may be absent and if not, it corresponds to the HTTP response code associated to the handled resource. The validation report is always present. Some properties are optional and added when requested using request parameters. See [Form data operations](#) for more information.

```
{  
  "validation": [  
    {  
      "level": "error",
```

```

        "message": "Field 'Value' is mandatory.",
        "blocksCommit": "never",
        "pathInRecord": "/value"
    },
    {
        "level": "error",
        "message": "Values between 'AD100' and 'AD200' (included) are
prohibited.",
        "blocksCommit": "onUserSubmit-checkModifiedValues",
        "pathInRecord": "/code"
    },
    {
        "level": "warning",
        "message": "Value 'AD150' is prohibited.",
        "blocksCommit": "never",
        "pathInRecord": "/code"
    }
]
}

```

- If the insert multiple record request was successful, the `code` JSON property may be absent and if not, it corresponds to the HTTP response code associated to the handled resource. The validation report is always present. Some properties are optional and added when requested using request parameters.


```

{
  "rows": [
    {
      "label": "My Lable1",
      "details": "http://.../root/table/1",
      "foreignKey": "1",
      "validation": [
        {
          "level": "error",
          "message": "Field 'Value' is mandatory.",
          "blocksCommit": "never",
          "pathInRecord": "/value"
        },
        {
          "level": "error",
          "message": "Values between 'AD100' and 'AD200' (included) are prohibited.",
          "blocksCommit": "onUserSubmit-checkModifiedValues",
          "pathInRecord": "/code"
        },
        {
          "level": "warning",

```

```

        "message": "Value 'AD150' is prohibited.",

        "blocksCommit": "never",

        "pathInRecord": "/code"

    }

]

},

{

    "label": "My Lable2",

    "details": "http://.../root/table/2",

    "foreignKey": "2",

    "validation": [

        {

            "level": "error",

            "message": "Field 'Value' is mandatory.",

            "blocksCommit": "never",

            "pathInRecord": "/value"

        },

        {

            "level": "error",

            "message": "Values between 'AD100' and 'AD200' (included) are
prohibited.",

            "blocksCommit": "onUserSubmit-checkModifiedValues",

```

```

        "pathInRecord": "/code"
    },
    {
        "level": "warning",
        "message": "Value 'AD150' is prohibited.",
        "blocksCommit": "never",
        "pathInRecord": "/code"
    }
]
}
]
}

```

- If the insert/update of a single record request or update field in table/dataset failed, the response body corresponds to a JSON `Exception handling` response.

```

{
    "code": 422,
    "errors": [
        {
            "level": "error",
            "userCode": "Validation",
            "message": "Field 'Category' is mandatory.",
            "blocksCommit": "never",

```

```

        "pathInRecord": "/category"
    },
    {
        "level": "error",
        "userCode": "Validation",
        "message": "Values between 'AD100' and 'AD200' (included) are
prohibited.",
        "blocksCommit": "onUserSubmit-checkModifiedValues",
        "pathInRecord": "/code"
    }
]
}

```

- If the insert multiple record request failed, the response body corresponds to a JSON `Exception handling` response.

```

{
    "code": 422,
    "errors": [
        {
            "level": "error",
            "rowIndex": 0,
            "userCode": "Validation",
            "message": "Values between 'AD100' and 'AD200' (included) are
prohibited.",

```

```

        "blocksCommit": "onUserSubmit-checkModifiedValues",

        "pathInRecord": "/code"
    },

    {

        "level": "error",

        "rowIndex": 0,

        "userCode": "Validation",

        "message": "Field 'Value Less Than' is mandatory.",

        "blocksCommit": "never",

        "pathInRecord": "/less_than_value"
    },

    {

        "level": "error",

        "rowIndex": 1,

        "userCode": "Validation",

        "message": "Values between 'AD100' and 'AD200' (included) are
prohibited.",

        "blocksCommit": "onUserSubmit-checkModifiedValues",

        "pathInRecord": "/code"
    },

    {

        "level": "error",

```

```

    "rowIndex": 1,

    "userCode": "Validation",

    "message": "Field 'Value Less Than' is mandatory.",

    "blocksCommit": "never",

    "pathInRecord": "/less_than_value"

  }

]

}

```

See also:

- [Form validation report](#)
- [Exception handling](#)

Form validation report

This report can be retrieve by using the [Form data operations](#). It provides every validation constraint on the requested resource.

Those constraints have the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>message</code>	<code>String</code>	Description of the constraint.	Yes
<code>blocksCommit</code>	<code>String</code>	Control policy of the constraint. The possible values are: <code>onInsertUpdateOrDelete</code> , <code>onUserSubmit</code> - <code>checkModifiedValues</code> ,	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		<code>never</code> . See Blocking and non-blocking constraints for more information.	
<code>level</code>	<code>String</code>	Severity level of the constraint. The possible values are: <code>info</code> , <code>warning</code> , <code>error</code> or <code>fatal</code> .	Yes
<code>pathInDataset</code>	<code>String</code>	Relative field path starting from the schema node.	No(**)
<code>pathInRecord</code>	<code>String</code>	Relative field path starting from the table node.	No (*)

(*) Only available for record and record field operations.

(**) Only available for dataset operations.

See also:

- [Form data category](#)

Content

Pagination

This feature allows to return a limited and parameterizable number of data. Pagination can be applied to data of the following types: records, association values, selection node values, selectors and dataspace. A context named `pagination` is always returned. This context allows browsing data similarly to the UI.

The pagination is always enabled.

See also:

- [Select operation](#)
- [Beta feature: Select dataspace or snapshots](#)

Detailed information related to this context can be found hereafter:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>firstPage</code>	<code>String</code> or <code>null</code> (*)	URL to access the first page.	Yes (**)
<code>previousPage</code>	<code>String</code> or <code>null</code> (*)	URL to access the previous page.	Yes (**)
<code>nextPage</code>	<code>String</code> or <code>null</code> (*)	URL to access the next page.	Yes
<code>lastPage</code>	<code>String</code> or <code>null</code> (*)	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.

Content of simple type

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE	OPENAPI
<code>xs:string</code> <code>xs:Name</code> <code>osd:text</code>	String (Unicode characters, cf. RFC4627)	"A text" "The escape of \"special character\" is preceded by a backslash." null	<code>string</code> <code>name</code> <code>text</code>	type: <code>string</code> format: <code>n/a</code>
<code>osd:html</code>	String (Unicode characters, cf. RFC4627)	" An HTML tag can thus be written without trouble "	<code>html</code>	type: <code>string</code> format: <code>html</code>
<code>osd:email</code>	String (Unicode characters, cf. RFC4627)	"employee@mycompany.com"	<code>email</code>	type: <code>string</code> format: <code>email</code>
<code>osd:locale</code>	String (Language tag, cf. RFC1766)	"en-US"	<code>locale</code>	type: <code>string</code> format: <code>locale</code>
<code>xs:string</code> (Foreign key)	String contains the value of the formatted foreign key.	<code>"0"</code> <code>"true 99"</code>	<code>foreignKey</code>	type: <code>string</code> format: <code>n/a</code>
<code>xs:boolean</code>	Boolean	true false	<code>boolean</code>	type: <code>boolean</code> format: <code>n/a</code>

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE	OPENAPI
		null		
<code>xs:decimal</code>	Number or null	-10.5 20.001 15 -1e-13	decimal	type: number format: double
<code>xs:date</code>	String with format: "yyyy-MM-dd"	"2015-04-13"	date	type: string format: date
<code>xs:time</code>	String with format: "HH:mm:ss" "HH:mm:ss.SSS"	"11:55:00" "11:55:00.000"	time	type: string format: date-time
<code>xs:dateTime</code>	String with format: "yyyy-MM-ddTHH:mm:ss" "yyyy-MM-ddTHH:mm:ss.SSS"	"2015-04-13T11:55:00" "2015-04-13T11:55:00.000"	dateTime	type: string format: date-time
<code>xs:anyURI</code>	String (Uniform Resource Identifier, cf. RFC3986)	"https://fr.wikipedia.org/wiki/René_Descartes"	anyURI	type: string format: uri
<code>xs:int</code>	Number or null	1596	int	type: integer

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE	OPENAPI
<code>xs:integer</code>				format: <code>int32</code>
<code>osd:resource</code>	<code>String</code> contains the resource formatted value.	"ebx-tutorial:ext-images:frontpages/Bach.jpg"	<code>resource</code>	type: <code>string</code> format: <code>n/a</code>
<code>osd:color</code>	<code>String</code> with format: "#[A-Fa-f0-9]{6}" contains the formatted value for the color.	"#F6E0E0"	<code>color</code>	type: <code>string</code> format: <code>n/a</code>
<code>osd:datasetName</code>	<code>String</code> with format: "[a-zA-Z_][a-zA-Z0-9_]*" and 64 characters max. contains the formatted value of the dataset name.	"ebx-tutorial"	<code>dataset</code>	type: <code>string</code> format: <code>n/a</code>
<code>osd:dataspaceKey</code>	<code>String</code> with format: <code>[BV][a-zA-Z0-9_:.\\-\\]+</code> and 33	"Bebx-tutorial"	<code>dataspace</code>	type: <code>string</code> format: <code>n/a</code>

XML SCHEMA	JSON FORMAT	EXAMPLES	META TYPE	OPENAPI
	characters max. contains the formatted key value of the dataspace.			

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:

- `count` contains an `int` of JSON `Number` type, that defines the number of processed records since the listed rows can not be exhaustive in several cases.
- `isPartialList` contains a `boolean` of JSON `Boolean` type, that is true if the listed rows are only a partial list of the whole processed records, false otherwise.
- `rows` contains a JSON `Array`, where each element corresponds to the result of a request element.
- `code` contains an `int` of 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 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 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 JSON `String` type, corresponding to the resource URL. This property is included if, and only if, the parameter `includeDetails` is set to `true`.

```
{
  "count": 2,
  "isPartialList": false,
```

```
"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](#)

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 JSON `Number` type, corresponds to the number of deleted records.
- `occultedCount` containing an integer of JSON `Number` type, corresponds to the number of occulted records.
- `inheritedCount` containing an integer of JSON `Number` type, corresponds to the number of inherited records.

```
{
  "deletedCount": 1,
  "inheritedCount": 0,
  "occultedCount": 0
}
```

See also:

- [Delete operation](#)

Technical data

Each returned record can contain its system records metadata under `ebx-metadata/system` properties corresponding to the technical data. System metadata are requestable and sortable.

To include records metadata, see `includeMetadata=system` parameter value for select operation.

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>uuid</code>	<code>String</code>	Universally Unique Identifier of the record	Yes
<code>creator</code>	<code>String</code>	Creation user's identifier. See	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		UserReference for more information.	
<code>creation_time</code>	<code>String</code>	Creation date	Yes
<code>updater</code>	<code>String</code>	Last update user's identifier. See UserReference for more information.	Yes
<code>update_time</code>	<code>String</code>	Last update date.	Yes

For extended format:

```
{
  ...,
  "ebx-metadata":{
    "content":{
      "system":{
        "content":{
          "uuid":{
            "content":"7FE03810-6A67-11ED-A892-00FF20216100"
          },
          "creator":{
            "content":"Uadmin"
          },
          "creation_time":{
            "content":"2022-11-22T14:13:47.793"
```

```
    },  
  
    "updater": {  
  
        "content": "Uadmin"  
  
    },  
  
    "update_time": {  
  
        "content": "2022-11-29T17:46:55.686"  
  
    }  
  
}  
  
}  
  
}
```

For compact format:

```
{
  ...,
  "ebx-metadata": {
    "system": {
      "uuid": "7FE03810-6A67-11ED-A892-00FF20216100",
      "creator": "Uadmin",
      "creation_time": "2022-11-22T14:13:47.793",
      "updater": "Uadmin",
    }
  }
}
```



```

      "update_time": "2022-11-29T17:46:55.686"
    }
  }
}

```

See also:

- [Optimistic locking](#)

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](#) and [insert](#), as well as

`ImportSpec.setByDelta` in the Java API for more information.

STATE IN THE JSON SOURCE DOCUMENT	BEHAVIOR
The property does not exist in the source document	If the <code>byDelta</code> mode is activated (default): • For the <code>update</code> operation, the field value remains unchanged. • For the <code>insert</code> operation, the behavior is the same as when the <code>byDelta</code> mode is disabled.

STATE IN THE JSON SOURCE DOCUMENT	BEHAVIOR
	If the <code>byDelta</code> mode is disabled through the RESTful operation parameter: The target field is set to one of the following values: • If the element defines a default value, the target field is set to that default value. • If the element is of a type other than a string or list, the target field value is set to <code>null</code>. • If the element is an aggregated list, the target field value is set to an empty list value. • 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. • If the element (simple or complex) is hidden in the data services, the target value remains unchanged. See also: ◦ Hiding a field in Data Services Note: The user performing the import must have the required permissions to create or change the target field value. Otherwise, the operation will be aborted.

STATE IN THE JSON SOURCE DOCUMENT	BEHAVIOR
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.

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.

Indexing and search strategy

[Filterable](#) and [sortable](#) values from the metamodel are limited to the default search strategy.

See also:

- [Search](#)

Others

- [Introduction](#)
- [Global structure](#)
 - [JSON Request body](#)
 - [JSON Response body](#)

Introduction

The [Extended](#) and [Compact](#) JSON formats can handle most of the use cases, however some operations and REST categories require specific formats like the followings.

Global structure

JSON Request body

The Request body is represented by a JSON `Object` whose content varies according to the operation and the 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](#)

- **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](#)

JSON Response body

The response body is represented by a JSON `Object` whose content depends on the operation and the category.

Admin category

The selection operation for this category only provides the requested values 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](#)

Auth category

The response body contains several properties directly placed in the root JSON `Object`.

- **Token creation**

Contains the token value and its type.

```
{
  "accessToken": "...",          // JSON String
  "tokenType": "..."          // JSON String
}
```

See also:

- [Create token operation](#)

Data category

Look up table views

The response body contains a `content` property holding a JSON `Array`, itself composed by JSON `Object` s with the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>details</code>	<code>String</code>	Corresponds to the view access URL.	Yes
<code>label</code>	<code>String</code>	View's label.	No
<code>viewPublication</code>	<code>String</code>	Published view's name. See also:	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		Table viewPublication parameter	
viewType	String	Enumeration whose value corresponds to one of the following: - SimpleTabular : Simple tabular view. - Hierarchy : Hierarchical view. - Tile : Tile view.	Yes

```
{
  "content": [
    {
      "details": "http://.../data/v1/Bebx-directory/ebx-directory/
directory/users?viewPublication=custom-directory",
      "label": "My custom directory view",
      "viewPublication": "custom-directory",
      "viewType": "SimpleTabular"
    },
    {
```

```

    ...
  }
]
}

```

See also:

- [Look up table views operation](#)

Beta feature: Dataspaces selection

The returned response body contains dataspace in a `rows` JSON `Array` property, where each inner JSON `Object` corresponds to a dataspace with the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>label</code>	<code>String</code>	Documentation label in the current locale.	No
<code>description</code>	<code>String</code>	Documentation description in the current locale.	No
<code>details</code>	<code>String</code>	Specifies the dataspace's REST resource URL.	Yes
<code>information</code>	<code>String</code>	Specifies the dataspace's information REST resource URL.	Yes

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>key</code>	<code>String</code>	Specifies the dataspace or snapshot formatted key. Format: <code>[BV][a-zA-Z0-9_:. \- \]{1,33}</code> and percent encoded afterward.	Yes
<code>isSelectAllowed</code>	<code>Boolean</code>	Specifies if the dataspace can be selected, according to the user's permissions.	Yes
<code>hasChildren</code>	<code>Boolean</code>	Specifies if the dataspace has children, according to the user's permissions. Note: Not applicable for snapshots.	Yes
<code>children</code>	<code>String</code>	Specifies the dataspace's children REST resource URL. If <code>hasChildren</code> property key value is <code>false</code> then the	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		returned value is <code>null</code>. Note: Not applicable for snapshots.	
<code>snapshots</code>	<code>String</code>	Specifies the dataspace's snapshots REST resource URL. Note: Not applicable for snapshots.	Yes

```
{
  "rows": [
    {
      "label": "Master Data - Reference",
      "description": "Reference dataspace in EBX.",
      "details": "http://.../data/v1/BReference",
      "information": "http://.../data/v1/BReference:information",
      "key": "BReference",
      "closed": false,
      "isSelectAllowed": true,
      "hasChildren": true,
```

```
    "children": "http://.../data/v1/BReference:children",

    "snapshots": "http://.../data/v1/BReference:snapshots"

  },

  {

    // An other dataspace

  }

],

"pagination": {

  "firstPage": null,

  "previousPage": null,

  "nextPage": null,

  "lastPage": null

}

}
```

See also:

- [Pagination](#)
- [Beta feature: Select dataspace or snapshots](#)

Beta feature: Dataspace information

The response body contains a `content` JSON `Object` property, holding the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>content</code>	<code>Object</code>	Corresponds to the dataspace or the snapshot information.	Yes
<code>key</code>	<code>String</code>	Specifies the dataspace or snapshot formatted key value. Format: <code>[BV][a-zA-Z0-9_:. \-\\]</code> + limited to 33 characters maximum.	Yes
<code>documentation</code>	<code>Array</code> of <code>JSON Object</code>	Corresponds to the localized documentation with a <code>JSON Object</code> by locale.	No
<code>locale</code>	<code>String</code>	Documentation locale (nested under the <code>documentation</code> property).	Yes
<code>label</code>	<code>String</code>	Documentation label (nested under the <code>documentation</code> property).	No
<code>description</code>	<code>String</code>	Documentation description (nested under the	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		<code>documentation</code> property).	
<code>closed</code>	<code>Boolean</code>	Specifies if the dataspace is closed.	Yes
<code>locked</code>	<code>Boolean</code>	Specifies if the dataspace is locked. Note: Not applicable for snapshots .	Yes
<code>parent</code>	<code>String</code>	Specifies the parent dataspace or snapshot formatted key value.	No
<code>administration</code>	<code>Boolean</code>	Specifies if the dataspace or the snapshot is an administration one.	Yes
<code>owner</code>	<code>String</code>	Specifies the owner profile.	No
<code>creator</code>	<code>String</code>	Specifies the creator profile.	Yes
<code>creationDate</code>	<code>Date</code>	Specifies the creation date.	Yes

```

{
  "content": {
    "key": "BReference",
    "documentation": [
      {
        "locale": "en-US",
        "label": "Master Data - Reference",
        "description": "Reference dataspace in EBX."
      },
      {
        "locale": "fr-FR",
        "label": "Données - Référence",
        "description": "Espace de données référence de EBX."
      }
    ],
    "closed": false,
    "locked": false,
    "parent": null,
    "administration": false,
    "relationalMode": false,
    "owner": "Badministrator",
    "creator": "Badministrator",
  }
}

```

```

    "creationDate": "2019-04-28T19:49:04.838"

  }

}

```

See also:

- [Beta feature: Select dataspace or snapshots](#)

Beta feature: Dataspace child or snapshot creation

The response body contains a `content` JSON `object` property, holding the following properties:

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>key</code>	<code>String</code>	Specifies the dataspace or snapshot formatted key value. Format: <code>[BV][a-zA-Z0-9_:. \-\\]</code> <code>+</code> limited to 33 characters maximum. The default value is a timestamp.	No
<code>owner</code>	<code>String</code>	Specifies the owner profile. Default value is <code>null</code>.	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
<code>documentation</code>	Array of JSON Object	Corresponds to the localized documentation with a JSON Object by locale. Default value is <code>null</code> .	No
<code>locale</code>	String	Documentation locale (nested under the <code>documentation</code> property).	Yes
<code>label</code>	String	Documentation label (nested under the <code>documentation</code> property). Default value is <code>null</code> .	No
<code>description</code>	String	Documentation description (nested under the <code>documentation</code> property). Default value is <code>null</code> .	No
<code>dataspaceKeyToCopyFrom</code>	String	Specifies the dataspace's formatted key value	No

JSON PROPERTY	JSON FORMAT	DESCRIPTION	REQUIRED
		from which to copy permissions. Note: Not applicable for snapshots creation.	

```
{
  "content": {
    "key": "BMyData",
    "owner": "Beveryone",
    "documentation": [
      {
        "locale": "en-US",
        "label": "My dataspace",
        "description": "This space contains my data"
      }
    ],
    "dataspaceKeyToCopyPermissionsFrom": "BReference"
  }
}
```

See also:

- Beta feature: Create a child dataspace or a snapshot

Health category

Health started and check responses

The response body contains a JSON String value corresponding to an `OK` value.

Health stopped response

The response body contains a JSON String value corresponding to an enumeration value.

Possible values: `OK` or `NO`

SQL in EBX®

- [Introduction](#)
- [Comparison operators](#)
- [Arithmetic operators and functions](#)
- [Logical operators](#)
- [String operators and functions](#)
- [Date and time functions](#)
- [EBX® SQL functions](#)

Introduction

- [Overview](#)
- [Mapping data model entities](#)
 - [Table \(in data model\)](#)
 - [Fields](#)
 - [Groups](#)
 - [Associations](#)
- [Mapping data types](#)
 - [Simple data types](#)
 - [Extended simple types defined by EBX®](#)

- List (multi-valued) types
- Complex types
- TableRef types
- Metadata fields
- System columns
- SQL syntax
 - Supported standard operators and functions
 - Escaping identifiers
 - Explain plan
- Limitations and performance guidelines

Overview

This documentation covers Structured Query Language (SQL) queries and expressions in EBX®. EBX® supports standard SQL queries to retrieve rows selected from one or more tables. Some EBX® SQL language features are extensions to the standard. Supported EBX® SQL syntax includes: table expressions (SELECT, FROM, WHERE, GROUP BY and HAVING clauses), DISTINCT, ORDER BY, LIMIT and OFFSET, combining queries (UNION), and WITH (except RECURSIVE modifier).

The goal of this API is to provide to developers the ability to retrieve data from EBX® using a well-known standard.

EBX® SQL is accessible through Java APIs, especially from the class `Query`. The queries also support parameters. See `Query.setParameter`.

Queries can involve a single dataset or multiple datasets, see `QueryBuilder` to know more about how to create a query.

Mapping data model entities

The following section provides a detailed explanation about the mapping of the EBX® concepts into SQL.

Table (in data model)

EBX® tables are mapped naturally to SQL tables. In the data model, there can be more than one EBX® table with the same name. This ambiguity can occur when tables are in groups. To remove the ambiguity, use the full path of the table surrounded by double quotes (for example,

`"my_group/my_table"` no longer conflicts with `"other_group/my_table"`). You can also use the entity name of the table, which is unique inside the data model. You can use the table name only if it does not collide with an entity name or another table name.

See also:

- [Entity mappings using the Data Model Assistant](#)

Fields

In SQL Standard, the structure of a table consists of one or more columns. Every element (including fields) whose parent is an EBX® table, is mapped to a column.

Groups

In SQL Standard, the structure of a table consists of one or more columns. Every element whose parent is an EBX® table is mapped to a column. This includes groups that are mapped to SQL columns as *SQL structure types*.

Associations

In SQL Standard, querying data among multiple tables is based on foreign keys and primary keys. These concepts in EBX® are similar to those in SQL. Therefore, joins between tables in SQL can also be done using EBX® foreign and primary keys.

Mapping data types

Handling data through SQL is highly dependent on its data type. For example, in predicates, columns can be compared only if they have the same SQL data type. The SQL data types are types according to the type in the data model.

Simple data types

Supported standard SQL data types

This table lists all of the simple types defined in the XML Schema that are supported by EBX®, along with their corresponding standard SQL types.

XML SCHEMA TYPE	SQL TYPE	JAVA TYPE	NOTES
<code>xs:string</code>	<code>VARCHAR</code>	<code>java.lang.String</code>	
<code>xs:boolean</code>	<code>BOOLEAN</code>	<code>java.lang.Boolean</code>	Values: TRUE, FALSE, UNKNOWN
<code>xs:decimal</code>	<code>DECIMAL</code>	<code>java.math.BigDecimal</code>	
<code>xs:dateTime</code>	<code>TIMESTAMP</code>	<code>java.util.Date</code>	
<code>xs:time</code>	<code>TIME</code>	<code>java.util.Date</code>	The date portion of the returned <code>Date</code> is always set to '1970/01/01'.
<code>xs:date</code>	<code>DATE</code>	<code>java.util.Date</code>	The time portion of the returned <code>Date</code> is always set to the beginning of the day (that is, '00:00:00').
<code>xs:anyURI</code>	<code>VARCHAR</code>	<code>java.net.URI</code>	
<code>xs:Name</code> (<code>xs:string</code> restriction)	<code>VARCHAR</code>	<code>java.lang.String</code>	
<code>xs:int</code>	<code>INT</code>	<code>java.lang.Integer</code>	

Extended simple types defined by EBX®

EBX® provides pre-defined simple data types. These types are defined by the internal schema `common-1.0.xsd`. Their definition is detailed in the section [Extended simple types defined by EBX®](#)

XML SCHEMA TYPE	SQL TYPE	JAVA CLASS
<code>osd:text</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.lang.String</code>
<code>osd:html</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.lang.String</code>
<code>osd:email</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.lang.String</code>
<code>osd:password</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.lang.String</code>
<code>osd:color</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.lang.String</code>
<code>osd:resource</code> (<code>xs:anyURI</code> restriction)	VARCHAR	internal class
<code>osd:locale</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.util.Locale</code>
<code>osd:dataspaceKey</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.lang.String</code>
<code>osd:datasetName</code> (<code>xs:string</code> restriction)	VARCHAR	<code>java.lang.String</code>

List (multi-valued) types

Lists are handled as SQL Arrays. Their corresponding Java class is `java.util.List`.

Complex types

Complex types are handled as [SQL Structured types](#). Their corresponding Java class is `Object[]`. This applies to foreign keys (see below) and [groups](#), because they are defined through complex types. Use the *dot operator* to access fields inside the SQL Structure types. For example, use `address.street` to access the field `street` of the field `address`, if it is a complex type. When you reference a sub-field of a complex type in a query, you must always use the table name or alias:

- `SELECT customer.address.street FROM customer`
- `SELECT c.address.street FROM customer c`

TableRef types

In EBX®, a table can have a primary key composed of multiple fields. Foreign keys are defined by a single field with the `osd:tableRef` declaration. The standard SQL syntax has been extended to extract the value of any targeted primary key field. In the [Extraction of foreign keys example](#), the following SQL expressions are valid:

- `tableA.fkb.id = 123`
- `YEAR(tableA.fkb.date) > 2018`

Note:

Even if the primary key is composed of only one field, the name of the field must be specified to access the value. For example, if the primary key is composed of a single `id`, `fkb.id` must be used to access the value, as in `tableA.fkb.id = 123`

A typical example of tableRef usage would be SQL joins:

```
SELECT * FROM employee JOIN department ON employee.fkDept.id = department.id
```

Or in case the referenced table has a composite primary key:

```
SELECT * FROM tableA JOIN tableB ON tableA.fkB.id1 = tableB.id1 AND  
tableA.fkB.id2 = tableB.id2
```

Metadata fields

EBX® SQL provides also access to some metadata information. These fields are not returned by a SQL statement, unless they are explicitly referenced. For example, "`SELECT * FROM ...`" does not return metadata fields, but "`SELECT metadataField FROM ...`" does.

NAME	DESCRIPTION	SQL TYPE	JAVA CLASS	EXAMPLES
<code>creator</code>	The <code>creator</code> of the record.	<code>VARCHAR</code>	<code>java.lang.String</code>	<code>SELECT t."ebx- metadata"."system"."creat FROM myTable</code>

NAME	DESCRIPTION	SQL TYPE	JAVA CLASS	EXAMPLES
				<pre> t SELECT t.* FROM myTable t WHERE t."ebx- metadata"."system"."creat = 'Uuser1' SELECT t.* FROM myTable t ORDER BY t."ebx- metadata"."system"."creat </pre>
creation_time	The creation time of the record.	TIMESTAMP	java.util.Date	<pre> SELECT t."ebx- metadata"."system"."creat FROM myTable t SELECT t.* FROM myTable t WHERE t."ebx- metadata"."system"."creat > TIMESTAMP'2001-06-10 14:57:20' SELECT t.* FROM myTable t ORDER BY t."ebx- metadata"."system"."creat </pre>
updater	The last user who modified the record.	VARCHAR	java.lang.String	<pre> SELECT t."ebx- metadata"."system"."updat FROM myTable t SELECT t.* FROM myTable t WHERE t."ebx- metadata"."system"."updat = 'Uuser1' </pre>

NAME	DESCRIPTION	SQL TYPE	JAVA CLASS	EXAMPLES
				<pre>SELECT t.* FROM myTable t ORDER BY t."ebx- metadata"."system"."update</pre>
update_time	The time of the last modification of the record.	TIMESTAMP	java.util.Date	<pre>SELECT t."ebx- metadata"."system"."update FROM myTable t SELECT t.* FROM myTable t WHERE t."ebx- metadata"."system"."update > TIMESTAMP '2001-06-10 14:57:20' SELECT t.* FROM myTable t ORDER BY t."ebx- metadata"."system"."update</pre>

System columns

Apart from the fields present in a table, EBX® SQL provides some extra system columns. These columns are not returned by a SQL statement, unless they are explicitly referenced. For example, "`SELECT * FROM ...`" does not return systems columns, but "`SELECT systemColumnName FROM ...`" does.

NAME	DESCRIPTION	SQL TYPE	JAVA CLASS	EXAMPLES
\$adaptation	The Adaptation representing the table record.	internal type	com.onwbp.adapt	<pre>SELECTdaptation t."\$adaptation" FROM myTable t WHERE t.value>100</pre>

NAME	DESCRIPTION	SQL TYPE	JAVA CLASS	EXAMPLES
\$pk	String representation of the primary key of the record. See also: • <code>PrimaryKey.syntax</code>	VARCHAR	<code>java.lang.String</code>	<pre>SELECT t.* FROM myTable1 t WHERE t."\$pk"='123'</pre> <pre>SELECT t.* FROM myTable2 t WHERE t."\$pk"='123 abc'</pre> <pre>SELECT t."\$pk" FROM myTable3 t WHERE t.value>100</pre> <pre>SELECT t.value FROM myTable3 t ORDER BY t."\$pk"</pre>

SQL syntax

Supported standard operators and functions

An operator is a reserved word or a character used primarily in a SQL statement's WHERE clause to perform operation(s), such as comparisons and arithmetic operations. These Operators are used to specify conditions in a SQL statement, and to serve as conjunctions for multiple conditions in a statement. EBX® supports most of the SQL standard operators and functions. Some functions and operators can have optional parameters: they are surrounded by square brackets in the documentation. Generally, there are five types of operators and functions:

- [Comparison operators and functions](#)

- Arithmetic operators and functions
- Logical operators
- String operators and functions
- Date and time functions

The following table lists all of the operators' associativity and precedence, highest to lowest.

OPERATOR	ASSOCIATIVITY
.	left
(array element)	left
+ - (unary plus, minus)	right
* / %	left
+ -	left
BETWEEN, IN, LIKE, CONTAINS, and so on	-
< > = <= >= <> !=	left
IS NULL, IS FALSE, IS NOT TRUE, and so on	-
NOT	right
AND	left
OR	left

Escaping identifiers

In the following cases, the identifier must be escaped by using double quotes:

- when using the absolute path to identify a table (for example, "/root/myTable").
- when the field to identify is a reserved word (for example, "user", "order").
- when referring to a system column with a table alias (for example, t."\$adaptation", t."\$pk").
- when referring to a metadata field (for example, t."ebx-metadata"."system"."creation_time").

The following example shows a query to illustrate all cases:

```
SELECT t."user", t."$pk" FROM "/root/myTable" t WHERE t."order" = 1
```

Explain plan

EBX® SQL supports `EXPLAIN PLAN FOR ...` syntax to get the plan information of a query.

The result is similar to `Query.explain`.

Example: `EXPLAIN PLAN FOR SELECT id FROM myTable`

Limitations and performance guidelines

- Certain internal join optimizations do not support RIGHT and FULL joins, so avoid these join types if possible.
- The maximum precision and scale for numeric or decimal values is 1000.
- Queries using `GROUP BY` and/or aggregate functions (`MIN` , `MAX` , and so on) are not optimized, except for `COUNT` , which can be optimized in some circumstances.
- Currently, `MIN` and `MAX` operators do not exploit internal indices. Instead, use the following equivalent queries, which are probably more efficient:

```
SELECT val FROM myTable ORDER BY val DESC NULLS LAST LIMIT 1
```

instead of `SELECT MAX(val) FROM myTable`

```
SELECT val FROM myTable ORDER BY val LIMIT 1
```

instead of `SELECT MIN(val) FROM myTable`

Comparison operators

The table below lists all the SQL arithmetic operators and functions supported by EBX®, along with their standard SQL syntax. Some functions may have optional parameters: they are surrounded by square brackets.

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
value1 = value2	Equals <code>SELECT 1 = 0 : false</code> <code>SELECT 4 = 4 : true</code>
value1 <> value2	Not equals <code>SELECT 1 <> 0 : true</code> <code>SELECT 4 <> 4 : false</code>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
value1 > value2	Greater than <code>SELECT 1 > 0 : true</code> <code>SELECT 4 > 4 : false</code>
value1 >= value2	Greater than or equal <code>SELECT 1 >= 0 : true</code> <code>SELECT 4 >= 4 : true</code>
value1 < value2	Lower than <code>SELECT 1 < 0 : false</code> <code>SELECT 4 < 4 : false</code>
value1 <= value2	Less than or equal <code>SELECT 1 <= 0 : false</code> <code>SELECT 4 <= 4 : true</code>
value IS NULL	Whether value is null <code>SELECT 1 IS NULL : false</code> <code>SELECT NULL IS NULL : true</code>
value IS NOT NULL	Whether value is not null <code>SELECT 1 IS NOT NULL : true</code> <code>SELECT NULL IS NOT NULL : false</code>
value1 IS DISTINCT FROM value2	Whether two values are not equal, treating null values as the same <code>SELECT 1 IS DISTINCT FROM 1 : false</code> <code>SELECT 1 IS DISTINCT FROM 4 : true</code> <code>SELECT 1 IS DISTINCT FROM NULL : true</code> <code>SELECT NULL IS DISTINCT FROM NULL : false</code>
value1 IS NOT DISTINCT FROM value2	Whether two values are equal, treating null values as the same <code>SELECT 1 IS NOT DISTINCT FROM 1 : true</code> <code>SELECT 1 IS NOT DISTINCT FROM 4 : false</code> <code>SELECT 1 IS NOT DISTINCT FROM NULL : false</code> <code>SELECT NULL IS NOT DISTINCT FROM NULL : true</code>
value1 BETWEEN value2 AND value3	Whether value1 is greater than or equal to value2 and less than or equal to value3 <code>SELECT 4 BETWEEN 3 AND 10 : true</code> <code>SELECT 1 BETWEEN 3 AND 10 : false</code>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
value1 NOT BETWEEN value2 AND value3	Whether value1 is greater than or equal to value2 and less than or equal to value3 <pre>SELECT 4 NOT BETWEEN 3 AND 10 : false SELECT 1 NOT BETWEEN 3 AND 10 : true</pre>
string1 LIKE string2	Whether string1 matches pattern string2. The wildcard '%' represent zero, one or multiple characters. To find if string1 starts with a sequence, use pattern 'sequence%'. The matching is case sensitive. To perform matching that is not case sensitive, use UPPER (or LOWER) on string1 and pattern. <pre>SELECT firstname FROM employee WHERE name LIKE 'S%' : John, Jennifer SELECT firstname FROM employee WHERE UPPER(name) LIKE 'SM%' : John SELECT firstname FROM employee WHERE name LIKE '_m%' : John</pre>
string1 NOT LIKE string2	Whether string1 does not matches pattern string2. The wildcard '%' represent zero, one or multiple characters. To find if string1 does not start with a sequence, use pattern 'sequence%'. The matching is case sensitive. To perform matching that is not case sensitive, use UPPER (or LOWER) on string1 and pattern. <pre>SELECT firstname FROM employee WHERE name NOT LIKE 'S%' : Maria SELECT firstname FROM employee WHERE UPPER(name) NOT LIKE 'SM%' : Maria, Jennifer SELECT firstname FROM employee WHERE name NOT LIKE '_m%' : Maria, Jennifer</pre>
value IN (value *)	Whether value is equal to a value in a list <pre>SELECT firstname FROM employee WHERE name IN ('Smith', 'Hamilton') : John, Maria</pre>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
value NOT IN (value *)	Whether value is not equal to every value in a list SELECT firstname FROM employee WHERE name NOT IN ('Smith', 'Hamilton') : Jennifer
value IN (sub-query)	Whether value is equal to a row returned by sub-query SELECT e.firstname FROM employee e WHERE e.department.id IN (SELECT d.id FROM department d WHERE d.name='IT') : John

Arithmetic operators and functions

The following table lists all of the SQL arithmetic operators and functions supported by EBX®, along with their standard SQL syntax. Some functions can have optional parameters: they are surrounded by square brackets.

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
+ numeric	Returns numeric SELECT +4 : 4
- numeric	Returns numeric SELECT -4 : -4
numeric1 + numeric2	Returns numeric1 plus numeric2 SELECT 18 + 4 : 22
numeric1 - numeric2	Returns numeric1 minus numeric2 SELECT 18 - 4 : 14
numeric1 * numeric2	Returns numeric1 multiplied by numeric2 SELECT 18 * 4 : 72
numeric1 / numeric2	Returns numeric1 divided by numeric2 SELECT 18 / 4 : 4
ABS(numeric)	Returns the absolute value of numeric SELECT ABS(-243.5) : 243.5
ACOS(numeric)	Returns the arc cosine of numeric SELECT ACOS(-1) : 3.141592653589793

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
ASIN(numeric)	Returns the arc sine of numeric <code>SELECT ABS(0.25) : 0.25</code>
ATAN(numeric)	Returns the arc tangent of numeric <code>SELECT ATAN(2.5) : 1.1902899496825317</code>
ATAN2(numeric, numeric)	Returns the arc tangent of the numeric coordinates <code>SELECT ATAN2(0.50, 1) : 0.4636476090008061</code>
CEIL(numeric)	Rounds numeric up, and returns the smallest number that is greater than or equal to numeric <code>SELECT CEIL(25.75) : 26</code>
CEILING(numeric)	Rounds numeric up, and returns the smallest number that is greater than or equal to numeric <code>SELECT CEILING(25.75) : 26</code>
COS(numeric)	Returns the cosine of numeric <code>SELECT COS(2) : -0.4161468365471424</code>
COT(numeric)	Returns the cotangent of numeric <code>SELECT COT(6) : -3.436353004180128</code>
DEGREES(numeric)	Converts numeric from radians to degrees <code>SELECT DEGREES(1.5) : 85.94366926962348</code>
EXP(numeric)	Returns e raised to the power of numeric <code>SELECT EXP(1) : 2.718281828459045</code>
FLOOR(numeric)	Rounds numeric down, and returns the largest number that is less than or equal to numeric <code>SELECT FLOOR(25.75) : 25</code>
LN(numeric)	Returns the natural logarithm (base e) of numeric <code>SELECT LN(2) : 0.6931471805599453</code>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
LOG10(numeric)	Returns the base-10 logarithm of numeric SELECT LOG10(2) : 0.3010299956639812
MOD(numeric1, numeric2)	Returns the remainder (modulus) of numeric1 divided by numeric2. The result is negative only if numeric1 is negative SELECT MOD(18, 4) : 2
POWER(numeric1, numeric)	Returns numeric1 raised to the power of numeric2 SELECT POWER(4, 2) : 16.0
RADIANS(numeric)	Converts numeric from degrees to radians SELECT RADIANS(180) : 3.141592653589793
RAND()	Returns a random double using numeric as the seed value if specified SELECT RAND(6) : 0.9891840064573959
RAND_INTEGER(numeric)	Generates a random integer between 0 and numeric - 1 inclusive, initializing the random number generator with seed if specified SELECT RAND_INTEGER(100, 5) : 0 SELECT RAND_INTEGER(6, 100) : 11
ROUND(numeric1, numeric2)	Rounds numeric1 to numeric2 places right to the decimal point SELECT ROUND(135.375, 2) : 135.38
SIGN(numeric)	Returns the signum of numeric SELECT SIGN(255.5) : 1
SIN(numeric)	Returns the sine of numeric SELECT SIN(2) : 0.9092974268256817
SQRT(numeric)	Returns the square root of numeric SELECT SQRT(64) : 8.0
TAN(numeric)	Returns the tangent of numeric SELECT TAN(1.75) : -5.52037992250933

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
TRUNCATE(numeric1)	Truncates numeric1 to 0 (if no numeric2 specified) places right to the decimal point <pre>SELECT TRUNCATE(135.375) : 135</pre> <pre>SELECT TRUNCATE(135.375, 2) : 135.37</pre>

Logical operators

The following table lists all of the SQL logical operators supported by EBX®, along with their standard SQL syntax.

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
boolean1 OR boolean2	Whether boolean1 is TRUE or boolean2 is TRUE <pre>SELECT TRUE OR TRUE : true</pre> <pre>SELECT TRUE OR FALSE : true</pre> <pre>SELECT FALSE OR TRUE : true</pre> <pre>SELECT FALSE OR FALSE : false</pre>
boolean1 AND boolean2	Whether boolean1 is TRUE and boolean2 is TRUE <pre>SELECT TRUE AND TRUE : true</pre> <pre>SELECT TRUE AND FALSE : false</pre> <pre>SELECT FALSE AND TRUE : false</pre> <pre>SELECT FALSE AND FALSE : false</pre>
NOT boolean	Whether boolean is not TRUE; returns UNKNOWN if boolean is UNKNOWN <pre>SELECT NOT TRUE : false</pre> <pre>SELECT NOT FALSE : true</pre> <pre>SELECT NOT UNKNOWN : null</pre>
boolean IS FALSE	Whether boolean is FALSE; returns FALSE if boolean is UNKNOWN <pre>SELECT TRUE IS FALSE : false</pre> <pre>SELECT FALSE IS FALSE : true</pre>
boolean IS NOT FALSE	Whether boolean is not FALSE; returns TRUE if boolean is UNKNOWN <pre>SELECT TRUE IS NOT FALSE : true</pre> <pre>SELECT FALSE IS NOT FALSE : false</pre>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
boolean IS TRUE	Whether boolean is TRUE; returns TRUE if boolean is UNKNOWN SELECT TRUE IS TRUE : true SELECT FALSE IS TRUE : false
boolean IS NOT TRUE	Whether boolean is not TRUE; returns FALSE if boolean is UNKNOWN SELECT TRUE IS NOT TRUE : false SELECT FALSE IS NOT TRUE : true
boolean IS UNKNOWN	Whether boolean is UNKNOWN SELECT TRUE IS UNKNOWN : false SELECT FALSE IS UNKNOWN : false
boolean IS NOT UNKNOWN	Whether boolean is not UNKNOWN SELECT TRUE IS NOT UNKNOWN : true SELECT FALSE IS NOT UNKNOWN : true

String operators and functions

The following table lists all of the SQL string operators and functions supported by EBX®, along with their standard SQL syntax. Some functions can have optional parameters: they are surrounded by square brackets.

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
string	
CHAR\ _LENGTH(string)	Returns the number of characters in a character string SELECT CHAR_LENGTH('Alfreds Futterkiste') : 19
CHARACTER\ _LENGTH(string)	As CHAR\ _LENGTH(string) SELECT CHARACTER_LENGTH('Alfreds Futterkiste') : 19
UPPER(string)	Returns a character string converted to upper case SELECT UPPER('SQL Tutorial is FUN!') : SQL TUTORIAL IS FUN!

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
POSITION(string1 IN string2)	Returns the position of the first occurrence of string1 in string2 starting at a given point if specified <pre>SELECT POSITION('is fun' IN 'Tutorial is FUN!') : 0</pre> <pre>SELECT POSITION('fun' IN 'Tutorial is FUN, very FUN!' FROM 17) : 0</pre>
TRIM({ BOTH	LEADING
OVERLAY(string1 PLACING string2 FROM integer)	Replaces a substring of string1 with string2 <pre>SELECT OVERLAY('Tutorial is very FUN!' PLACING 'VERY' FROM 13) :</pre> <pre>Tutorial is VERY FUN!</pre> <pre>SELECT OVERLAY('Tutorial is very FUN!' PLACING 'VERY' FROM 13 FOR 4) :</pre> <pre>Tutorial is VERY FUN!</pre>
SUBSTRING(string FROM integer)	Returns a substring of a character string starting at a given point <pre>SELECT SUBSTRING('Tutorial is very FUN!' FROM 13) : very FUN!</pre> <pre>SELECT SUBSTRING('Tutorial is very FUN!' FROM 13 FOR 4) : very</pre>

Date and time functions

The following table lists all of the SQL date and time functions supported by EBX®, along with their standard SQL syntax. Some functions can have optional parameters: they are surrounded by square brackets.

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
CURRENT_TIME	Returns the current time in the session time zone, in a value of datatype TIMESTAMP WITH TIME ZONE
CURRENT_DATE	Returns the current date in the session time zone, in a value of datatype DATE

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
CURRENT_TIMESTAMP	Returns the current date and time in the session time zone, in a value of datatype TIMESTAMP WITH TIME ZONE
YEAR(date)	Extracts and returns the value of the year from a datetime value expression. Returns an integer. <pre>SELECT YEAR(TIMESTAMP '1971-07-20 09:34:21') : 1971</pre> <pre>SELECT YEAR(DATE '1968-07-20') : 1968</pre> <pre>SELECT YEAR(hiringDate) FROM employee WHERE name='Smith' : 2015</pre> <pre>SELECT YEAR(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 2017</pre>
QUARTER(date)	Extracts and returns the value of the quarter from a datetime value expression. Returns an integer between 1 and 4. <pre>SELECT QUARTER(TIMESTAMP '1971-07-20 09:34:21') : 3</pre> <pre>SELECT QUARTER(hiringDate) FROM employee WHERE name='Smith' : 4</pre> <pre>SELECT QUARTER(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 2</pre>
QUARTER(date)	Extracts and returns the value of the month from a datetime value expression. Returns an integer between 1 and 12. <pre>SELECT MONTH(TIMESTAMP '1971-07-20 09:34:21') : 7</pre> <pre>SELECT MONTH(hiringDate) FROM employee WHERE name='Smith' : 10</pre> <pre>SELECT MONTH(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 6</pre>
WEEK(date)	Extracts and returns the value of the week from a datetime value expression. Returns an integer between 1 and 53. <pre>SELECT WEEK(TIMESTAMP '1971-07-20 09:34:21') : 29</pre> <pre>SELECT WEEK(hiringDate) FROM employee WHERE</pre>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
	<pre>name='Smith' : 42 SELECT WEEK(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 26</pre>
DAYOFYEAR(date)	Extracts and returns the value of the day of year from a datetime value expression. Returns an integer between 1 and 366. <pre>SELECT DAYOFYEAR(TIMESTAMP '1971-07-20 09:34:21') : 201 SELECT DAYOFYEAR(hiringDate) FROM employee WHERE name='Smith' : 287 SELECT DAYOFYEAR(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 179</pre>
DAYOFMONTH(date)	Extracts and returns the value of the day of month from a datetime value expression. Returns an integer between 1 and 31. <pre>SELECT DAYOFMONTH(TIMESTAMP '1971-07-20 09:34:21') : 20 SELECT DAYOFMONTH(hiringDate) FROM employee WHERE name='Smith' : 14 SELECT DAYOFMONTH(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 28</pre>
DAYOFWEEK(date)	Extracts and returns the value of the day of week from a datetime value expression. Returns an integer between 1 and 7. <pre>SELECT DAYOFWEEK(TIMESTAMP '1971-07-20 09:34:21') : 3 SELECT DAYOFWEEK(hiringDate) FROM employee WHERE name='Smith' : 4 SELECT DAYOFWEEK(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 4</pre>
HOUR(date)	Extracts and returns the value of the hours from a datetime value expression. Returns an integer between 0 and 23. <pre>SELECT HOUR(TIMESTAMP '1971-07-20 09:34:21') : 9 SELECT HOUR(TIME '10:34:21') : 10 SELECT HOUR(workStart) FROM employee WHERE</pre>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
	<pre>name='Smith' : 8 SELECT HOUR(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 16</pre>
MINUTE(date)	Extracts and returns the value of the minutes from a datetime value expression. Returns an integer between 0 and 59. <pre>SELECT MINUTE(TIMESTAMP '1971-07-20 09:34:21') : 34 SELECT MINUTE(TIME '10:35:21') : 35 SELECT MINUTE(workStart) FROM employee WHERE name='Smith' : 0 SELECT MINUTE(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 6</pre>
SECOND(date)	Extracts and returns the value of the seconds from a datetime value expression. Returns an integer between 0 and 59. <pre>SELECT HOUR(TIMESTAMP '1969-07-20 09:34:21') : 9 SELECT SECOND(TIME '10:34:22') : 22 SELECT SECOND(workStart) FROM employee WHERE name='Smith' : 0 SELECT SECOND(lastProfileUpdateTime) FROM employee WHERE name='Smith' : 12</pre>

EBX® SQL functions

The following table lists all of the EBX® built-in SQL functions, along with their syntax. Some functions can have optional parameters: they are surrounded by square brackets.

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
FK_AS_STRING('referencedDatasetName', 'referencedTableName', tableRefValue)	Returns the string representation of a tableRef value. See QueryBuilder to know more about dataset registered names <pre>SELECT e.name FROM employee e WHERE FK_AS_STRING('_public', '/root/ department', e.department) = '1' : Smith SELECT FK_AS_STRING('_public',</pre>

OPERATOR SYNTAX	DESCRIPTION AND EXAMPLE(S)
	<pre> '/root/department', e.department) FROM employee e WHERE e.name = 'Smith' : 1 </pre>

REST Toolkit

- Introduction
- Application definitions
- Service and operation definitions
 - URL and sample
- Serialization of a table record
 - Built-in serializers
 - Custom serializers
- Authentication and lookup mechanism
- REST authentication and permissions
- URI builders
 - URI builders for built-in
 - URI builders for resources
 - URI builders for server and application
- Exception handling
- Monitoring
- Packaging and registration
- OpenAPI
 - Overview
 - Activation
 - Configuration
 - Operations
 - Enhance documentation

- [Permissions](#)
- [Visual documentation](#)
- [Limitations](#)

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](#).

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 built URIs without worrying about the configuration (for example [reverse-proxy](#) or [REST forward](#) 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](#)

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` 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)` or `ApplicationConfigurator.register(java.lang.Object)` 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` 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()));
    }
}

```

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_URL_ENCODED_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).v
        alue();
    }
}

```

```

    }

    /**
     * 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()

            .stream()

            .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;

    }

}

```

Serialization of a table record

Built-in serializers

Default JSON serializers and deserializers are provided to handle table records when declared in DTOs as `ContentHolder`s. Both [extended](#) and [compact](#) JSON formats of record are supported.

```

/**

 * DTO for a singer.

 */

```

```

public final class SingerDTO

{

    @Table(

        dataModel = "urn:ebx:module:tracks-module:/WEB-INF/ebx/schemas/
tracks.xsd",

        tablePath = "/root/Singers")

    public ContentHolder content;

}

```

A same DTO can be used for serialization and deserialization. In case of serialization, a `ContentHolderForInput` instance will be automatically created and filled with the proper data. Afterwards, this instance will be able to copy its data into a `ValueContextForUpdate`. To deserialize a table records, a `ContentHolderForOutput` must be created in the REST operation JAVA method and returned. The provided `Adaptation` data will then be transformed into a valid piece of JSON and placed into the HTTP response body.

```

/**

 * The REST Toolkit Singer service v1.

 */

@Produces(MediaType.APPLICATION_JSON)

@Consumes(MediaType.APPLICATION_JSON)

@Path("/singer/v1")

@Documentation("Singer service")

public final class SingerService

{

    ...

```

```

/**

 * Selects a singer by id.

 */

@GET

@Path("/singers/{id}")

@Documentation("Selects a singer by id")

public SingerDTO handleSelectSingerById(final @PathParam("id") Integer
id)

{

    // find the singer adaptation by id

    final Adaptation singerRecord = ... ;

    final SingerDTO singerDTO = new SingerDTO();

    singerDTO.content =
ContentHolderForOutput.createForRecord(singerRecord);

    return singerDTO;

}

/**

 * Inserts one singer.

 */

@POST

```

```

@Path("/singers")

@Documentation("Inserts one singer")

public void handleInsertOneSinger(final SingerDTO aSingerDTO)

{

    final ProgrammaticService svc = ... ;

    final AdaptationTable singersTable = ... ;

    final ProcedureResult procedureResult = svc.execute(

        (aContext) -> {

            final ValueContextForUpdate createContext =
aContext.getContextForNewOccurrence(singersTable); ;

aSingerDTO.content.asContentHolderForInput().copyTo(createContext);

            aContext.doCreateOccurrence(createContext,
singersTable);

        });

    if (procedureResult.hasFailed())

        throw new UnprocessableEntityException(

            procedureResult.getException().getLocalizedMessage());

}

/**

 * updates one singer.

```

```

    */

    @PUT

    @Path("/singers/{id}")

    @Documentation("updates one singer")

    public void handleUpdateOneSinger(@PathParam("id") final Integer id,
    final SingerDTO aSingerDTO)

    {

        final ProgrammaticService svc = ... ;

        final AdaptationTable singersTable = ... ;

        final ProcedureResult procedureResult = svc.execute(

            (aContext) -> {

                // find the singer adaptation by id

                final Adaptation singerRecord = ... ;

                if (singerRecord == null)

                    throw new NotFoundException("Singer with the id
["+ id + "] has not been found.");

                final ValueContextForUpdate createContext =
aContext.getContext(singerRecord.getAdaptationName()); ;

                aSingerDTO.content.asContentHolderForInput().copyTo(createContext);

                aContext.doModifyContent(singerRecord, createContext);

            });

```



```

        if (procedureResult.hasFailed()){

            final Exception ex = procedureResult.getException();

            final Throwable cause = ex.getCause();

            if(cause instanceof NotFoundException)

                throw (NotFoundException) cause;

            throw new
UnprocessableEntityException(ex.getLocalizedMessage());

        }

    }

}

```

The default JSON format of the responses, only composed of business data fields, is called *compact*.

```

{

    "singer":{

        "firstname":"Frank",

        "lastname":"Sinatra"

    }

}

```

To add technical fields or metadata, the `ExtendedOutput` annotation must be placed over the `ContentHolder` field. The annotation must declare every wished options or the *ALL* one.

```

/**
 * DTO for a singer with technical fields.
 */

public final class SingerWithTechnicalsDTO
{
    @Table(
        dataModel = "urn:ebx:module:tracks-module:/WEB-INF/ebx/schemas/
tracks.xsd",
        tablePath = "/root/Singers")
    @ExtendedOutput({Include.LABEL, Include.TECHNICALS, Include.CONTENT})
    public ContentHolder content;
}

```

```

{
  "singer":{
    "label": "23",
    "creationDate": "2018-09-04T15:35:10.706",
    "creationUser": "user",
    "lastUpdateDate": "2018-10-02T17:05:47.090",
    "lastUpdateUser": "user",
    "content":{
      "firstname":{
        "content":"Frank",

```

```

        "label": "First Name"
    },
    "lastname": {
        "content": "Sinatra",
        "label": "Last Name"
    }
}
}
}
}

```

To use the extended JSON format for input HTTP requests the `ExtendedInput` annotation must be placed over the `ContentHolder` field.

```

/**
 * DTO for a singer using the extended JSON format.
 */
public final class SingerExtendedInputDTO
{
    @Table(
        dataModel = "urn:ebx:module:tracks-module:/WEB-INF/ebx/schemas/
tracks.xsd",
        tablePath = "/root/Singers")
    @ExtendedInput
    public ContentHolder content;
}

```

}

Custom serializers

Since TIBCO EBX® is based on JSON-B ([JSR-367](#)), custom serializers and deserializers can be defined through [JsonbTypeSerializer](#) and [JsonbTypeDeserializer](#) annotations.

```
/**
 * DTO for a vinyl.
 */
public final class VinylDTO
{
    @JsonbTypeSerializer(CustomTrackDtoSerializer.class)
    @JsonbTypeDeserializer(CustomTrackDtoDeserializer.class)
    public TrackDTO track;
}
```

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 `AnonymousAccessEnabled`. See [REST authentication and permissions](#) for more information.

See also:

- [Authentication](#)
- [Lookup mechanism](#)

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 `AnonymousAccessEnabled`.

Some methods may need to be restricted to given profiles. These methods may be annotated by `Authorization` to specify an authorization rule. An authorization rule is a Java class that implements the `AuthorizationRule` 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()

    {

        ...

    }

}

```

URI builders

REST Toolkit provides an utility interface `URIInfoUtility` to generate URIs. An instance of this interface is accessible through the injectable built-in object `SessionContext`.

URI builders for built-in

Several URI builders interfaces have been designed to ease built-in RESTful services URI build. Each interface constitute an aggregation of methods related to the same functional concept. This division allows development of modular and generic algorithms. Some of these interfaces are themselves aggregation of other ones, leading to intuitive use of the builders since only consistent combination of method calls are allowed. For example, a URI builder configured for the REST `hierarchy` category will not allow calls to record URI build methods, since record access are part of the REST `data` category concept.

Moreover, these URI builders are preconfigured according to:

- the data from the incoming HTTP request which constitute the base configuration,
- the `Module public path prefix` which prepends the EBX® module segment, if it exists,
- the `data services lineage setting` which overrides the whole previous configuration, if it exists,
- The `REST forward` mode which overrides partially the previous configuration since only a piece of the URI's path is kept, if it is activated.

See `URIBuilderForBuiltin` and `CategoryURIBuilder` for more information.

URI builders for resources

URI builders to generate resource access URI are also available. Since resources are seen as only one functional concept, every methods have been defined in the same interface.

Like the URI builder for built-in data services, these ones are already preconfigured according to:

- the data from the incoming HTTP request which constitute the base configuration,
- the `external resources settings` which overrides all or part of the previous configuration, if they exist,
- the `Module public path prefix` which prepends the EBX® module segment, if it exists.

URI builders for server and application

When URIs to services or resources not covered by the previous builders must be build, the default preconfigured URI builders should be used. There are two default URI builders which only differ in their ending path segment. The first one end its path at the server's segment (just before the EBX®

module segment) and the second at the REST application's last segment (defined in the `@ApplicationPath` annotation).

These default URI builders are already preconfigured according to:

- the data from the incoming HTTP request which constitute the base configuration,
- the [Module public path prefix](#) which prepends the EBX® module segment, if it exists.

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 `UnprocessableEntityException` that adds support for the HTTP `422` (*Unprocessable entity*) code.

Plain Java exceptions are mapped to the HTTP status code `500` (*Internal server error*).

```
{
    "errorCode": 999,           // JSON Number, EBX® error code used mainly
    for
                                // HTTP error 422 (optional)

    "localizedMessage": "...", // JSON String, localized message

    "messages": [              // JSON Array of String, Internal messages
    useful when debugging

                                // (optional).

                                // Usually not displayed to the user and not
    localized.

    "Message 1", "Message 2"

    ]

}
```


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](#)
- [Configuring the EBX® logs](#)

Some additional properties are available to configure the log messages. See [Configuring REST toolkit services](#) for further information.

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](#) for more information.

The registration of a REST Toolkit application is integrated into the EBX® module registration process. The registration class must extend `ModuleRegistrationListener`, declare the Servlet 3.0 annotation `@WebListener` and override the `handleContextInitialized` method.

See [Module registration](#) 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);
    }
}

```

OpenAPI

Overview

The OpenAPI documents, generated through the TIBCO EBX® native integration, comply with the [OpenAPI V3 specification](#) . By structuring and describing the available REST resources of a REST Toolkit application and associated the operations, These documents facilitate their development and consumption. However, only the JSON format is currently supported.

Activation

To activate the OpenAPI documents generation:

- The `OpenApiApplicationPath` annotation must be set on the REST Toolkit application class to document. This annotation requires the definition of a unique path used to configure the OpenAPI operations' base URI.

Note:

The OpenAPI application path will be followed by the generator version (i.e. `/open-api-application-path/v1`). It is recommended to not use paths like `/open-api` as there will be confusion between the generator version and the open-api specification version.

- The `OpenApiResource` annotation must be set on the REST Toolkit class resources to consider during the documentation process.

Configuration

The parameter `ebx.restservices.openapi.unpublished.applications` can be used to not publish the OpenAPI endpoints on a specific environment.

```
#####

# Comma-separated list of canonical application class names

# whose OpenAPI descriptions will not be published.

#####

ebx.restservices.openapi.unpublished.applications=com.example.RestApplication
```

The parameter `ebx.restservices.openapi.style.short.name` can be used to simplify the names of the generated schemas.

```
#####

# If true, only the simple name of the DTO classes will be used to name
the OpenAPI schemas.

# It also removes the DTO and Dto suffixes at the end of the class name.

# Default value is false.

# Note that this may cause conflicts between DTOs with the same name under
different packages.

# It is not recommended, when set to true, to use the underscore character
```

```
'_' in @Schema

# annotations since it cause calculation errors.

#####

ebx.restservices.openapi.style.short.name=true
```

Operations

The OpenAPI operations use the GET or POST HTTP method to generate the JSON document of an entire application or a single service.

The URL format for application is:

```
http[s]://<host>[:<port>]/<open-api-path>/v1
```

The URL format for a single service is:

```
http[s]://<host>[:<port>]/<open-api-application-path>/v1/{resourcePath:
[&#94;:]*}
```

Where:

- `<open-api-application-path>` corresponds to the unique path mentioned in the `OpenApiApplicationPath` annotation.
- `{resourcePath: [^:]*}` corresponds to the REST resource's path to document.

HTTP codes

HTTP CODE	DESCRIPTION
200 (OK)	The document has been successfully generated.
401 (Unauthorized)	Authentication has failed.

HTTP CODE	DESCRIPTION
403 (<i>Forbidden</i>)	The specified resource read permission has been denied to the current user.
404 (<i>Not found</i>)	The resource, specified in the URL, cannot be found.

Enhance documentation

MicroProfile OpenAPI

In addition to the JAX-RS annotations that will be processed, by default, to generate the OpenAPI documents, TIBCO EBX® supports the [Microprofile OpenAPI annotations](#). These annotations are used to enhance the description without needing to rewrite portions of the OpenAPI document that are already covered by the REST Toolkit framework. The annotations can be used on both application and resource classes depending on their type.

Note:

If not explicitly defined on a resource class, the [OpenAPIDefinition](#) annotation will be inherited from its REST Toolkit application class. This annotation describes the root document object of the OpenAPI document.

Serialization of a table record

There are further actions to take when a `ContentHolder` is used in a DTO to [serialize a table record](#).

- The `Table` annotation must be set on the field.
- The [Schema](#) annotation must tag the field for processing.

Note:

The `ref` value of the [Schema](#) annotation must be non empty and unique in the whole application scope.

```

/**
 * DTO for a singer.
 */

public final class SingerDTO

{

    @Table(

        dataModel = "urn:ebx:module:tracks-module:/WEB-INF/ebx/
schemas/tracks.xsd",

        tablePath = "/root/Singers")

    @Schema(ref = "singer")

    public ContentHolder content;

}

```

Note:

Note that due to a limitation the metadata field is not added in this case.

Permissions

In development run mode, anonymous access to OpenAPI endpoints is granted to any user. In production and integration run modes, the OpenAPI access permissions are inherited from the ones of the [underlying resources](#).

See also:

- [EBX® run mode](#)

- The resource's documentation will only contain the operations accessible by the current user.
- A permission denied will be returned if the current user can not access any operation.
- The REST Toolkit application's documentation will only contain the allowed scopes of resources.
- A permission denied will be returned if the current user can not access any resource.

Visual documentation

In the [development run mode](#), the [Swagger UI](#) is available through the UI tab Administration > Technical configuration > Modules and data models. It allows to view and interact with the documented resources of the corresponding REST Toolkit application. Visual documentation facilitates back-end implementation and client-side consumption.

The Swagger UI is also directly available through the following URLs:

URL format for application OpenAPI visualisation:

```
http[s]://<host>[:<port>]/<open-api-application-path>/v1/ui
```

URL format for a single service OpenAPI visualisation:

```
http[s]://<host>[:<port>]/<open-api-application-path>/v1/ui/{resourcePath:
[&#94;:]*}
```

Where:

- `<open-api-application-path>` corresponds to the unique path mentioned in the `OpenApiApplicationPath` annotation.
- `{resourcePath: [^:]*}` corresponds to the REST resource's path to document.

Limitations

OpenAPI	The metadata field is not added on table record serialization.

See [Serialization of a table record](#) for more information.

DTO classes using Java generics that are not lists or maps are not supported.

EBX® Scripting

- [Record permission](#)
- [Function field](#)
- [Function field API](#)
 - [Unit summary](#)
 - [Unit default](#)
 - [Unit core.complex](#)
 - [Unit core.date](#)
 - [Unit core.datetime](#)
 - [Unit core.list](#)
 - [Unit core.locale](#)
 - [Unit core.log](#)
 - [Unit core.math](#)
 - [Unit core.resource](#)
 - [Unit core.string](#)
 - [Unit core.time](#)
 - [Unit core.uri](#)

Record permission

- [Introduction](#)
- [Lexical structure](#)
 - [Introduction](#)
 - [Character set](#)
 - [Character case sensitivity](#)
 - [Comments](#)

- Keywords
- Identifiers
 - Unquoted identifier
 - Quoted identifiers
- Types
 - Supported types
 - Literals
- Field access
 - Simple fields
 - Foreign table fields
 - List (multi-valued) field access
 - Associations
- Operators
- Null value management
 - Arithmetic operators
 - Comparison operators
 - Boolean operators
 - Index expressions
 - Functions
- If statement
 - "If then" statement
 - "If then else" statement
 - "Then" or "else" body statement
- Return statement
- Context
 - Introduction
 - Dataspace

- Dataset
- Session
- Functions
 - Roles
 - String matching functions
 - Aggregate functions
 - Miscellaneous functions

Introduction

Use the record permission DSL (Domain Specific Language) to specify access rules on records of a given table.

The main goals of the record permission DSL are to not require Java coding, and to be easy to use by people without deep programming knowledge.

You can specify permission on any table using a script. The script consists of a sequence of **if then else** and **return** statements that indicate the permission for a record.

You can edit the script using the [Data Model Assistant \(DMA\)](#) .

Lexical structure

Introduction

The script has following structure:

```
begin

    <statement 1>

    <statement 2>

    ...

    <last statement>

end
```

All statements except the last one must be an "if then" or an "if then else" statement.

The last statement can be an "if then" , "if then else" or "return" statement.

Example:

```
if isMember(administrator) then

    // Administrator has no restrictions.

    return readWrite;

if isMember('french-team') and record.Country='F' then

    //Members of 'french-team' can view and modify data for France.

    return readWrite;

if isMember('us-team') and record.Country='US' then

    //Members of 'us-team' can view and modify data for US.

    return readWrite;

// This statement is not actually needed as 'hidden' is the default
permission.

return hidden;
```

Character set

The Unicode character set is supported.

Character case sensitivity

The DSL is case-sensitive.

Comments

A single line comment extends from `//` to the end of the current line:

```
// This is a comment

if record.LastName = 'Doe' then // This is another comment.

    return readOnly;
```

A multi-line comment extends from `/**/` and ends with `*/`:

```
/* This is a sample of a multi-line

    comment */

if record.isActive then

    return readWrite;
```

Keywords

There are two types of keyword:

- **Reserved keywords** are:
`if` , `then` , `else` , `begin` , `end` , `return` , `null` , `and` , `or` , `not` , `true` , `false` .
- **Unreserved keywords** are all other keywords defined by the DSL.

The main difference between the two types of keywords is that unreserved ones, but not reserved ones, can be used as plain (unquoted) identifiers.

Identifiers

Unquoted identifier

An **unquoted identifier** is an unlimited-length sequence of letters, digits or underscore (`\_`). The first character must be a letter or an underscore.

Valid letters are **a to z** and **A to Z** . Valid digits are **0 to 9** .

An unquoted identifier cannot be equal to a reserved keyword.

Quoted identifiers

A quoted identifier is an unlimited length of any Unicode character except a double quote (").

Quoted identifiers **must** be used surrounded by double quotes.

An unquoted identifier can be used surrounded by double quotes. This means that identifier "a_name" is equal to a_name.

Quoted identifiers can be reserved keywords.

Types

Supported types

The following types are supported:

TYPE	EBX® CORRESPONDING TYPES
Boolean	xs:boolean
Decimal	xs:decimal xs:int xs:integer
String	xs:string xs:anyURI xs:Name osd:text osd:html osd:email osd:password osd:color osd:resource osd:locale osd:dataspaceKey osd:datasetName
Timestamp	xs:dateTime
Date	xs:date
Time	xs:time

Literals

String literal

String literals can be any sequence of Unicode characters surrounded by single quotes. The following table displays characters that need to be replaced by an escape sequence:

CHARACTER	ESCAPE SEQUENCE
Tab	\t
Backspace	\b
New line	\n
Carriage return	\r
Form feed	\f
Single quote	\'
Backslash	\\

A character can be specified by a Unicode escape sequence that has the format \uXXXX with XXXX the hexadecimal code of the Unicode character.

Examples

VALUE	SYNTAX
O'Harra	'O\''Harra'
Noël	'No\u00EBI'
été	'\u00e9\u00E9'

Note

An invalid escape or Unicode sequence generates an error at compile time.

Decimal literal

The following decimal formats are supported:

FORMAT	EXAMPLES
Integer	546 -67
Floating point	54.987 -433.876 0.00054 -0.0032

FORMAT	EXAMPLES
Exponent notation	34.654e-5 -45E+65 1.543e23

Timestamp literal

Timestamp literals have the format **dt(yyyy-MM-dd hh:mm:ss.sss)** .

Seconds are optional. When seconds are not specified, 0 is assumed. Seconds can have fractions up to millisecond precision.

The dates that are not valid in the Gregorian calendar generate an error at compile time.

Examples:

```
dt(2010-01-02 00:00:00.000)
```

```
dt(2019-2-3 12:56:7)
```

```
dt(2019-2-3 12:56:7.5)
```

```
dt(2019-5-7 1:6)
```

Date literal

Date literals have format **d(yyyy-MM-dd)** .

The dates that are not valid in the Gregorian calendar generate an error at compile time.

Examples:

```
d(2010-01-02)
```

```
d(2019-2-3)
```

```
dt(2019-5-7)
```

Time literal

Time literals have the format **t(hh:mm:ss.sss)** .

Seconds are optional. When seconds are not specified, 0 is assumed. Seconds can have fractions up to millisecond precision.

Invalid times generate an error at compile time.

Examples:

```
t(00:00:00)

t(12:56:7)

t(12:56:7.5)

t(1:6)
```

Boolean literal

Boolean literals are **true** and **false**.

Field access

Simple fields

Only access to optimized/indexed table field is supported. Dot notation is used to access tables fields. For example, a condition of current table fields whose path is **.OfficeAddress/City** would be:

```
if record.OfficeAddress.City = 'Paris' then

    return readWrite;
```

The alias **record** always refers to the current record. Depending on context, other aliases might be available.

Each step (parts separated by a dot) is an **identifier**. This means that the following quoted notation can be used for any step:

```
if record."OfficeAddress".City = 'Paris' then

    return readWrite;
```


This is useful for steps equal to a reserved keyword or using characters, such as the minus character (-) or dot (.), which are not compatible with unquoted identifiers.

At runtime, any step can evaluate to **null** . In this case, the full field expression evaluates to **null** .

Foreign table fields

You can access foreign tables by "following" foreign keys using dot notation.

In the following example, the field **Supervisor** is a foreign key:

```
if record.Supervisor.Name = 'John Doe' then

    return readWrite;
```

There can be multiple levels of foreign keys, such as in the following example:

```
if record.Supervisor.Supervisor.Supervisor.Name = 'John Doe' then

    return readOnly;
```

List (multi-valued) field access

Multi-valued fields are not supported.

Associations

Aggregate functions can take as input expressions based on an association. In the following example, the field **ManagedUsers** is an association:

```
// All users that manage at least 2 persons have read write access.

if count(record.ManagedUsers[]) >= 2 then

    return readWrite;
```

You can apply a filter on an association. You must use an alias to access fields from the association. An association alias is declared using : . In the following example, the alias is **u1** :

```
// All users that manage at least one person whose office is in Briton has
read only access.

if exists(record.ManagedUsers:u1[u1.OfficeAddress.City='Briton']) then

    return readOnly;
```

A filter on an association can reference a field of the current record:

```
// All users that manage at least one person whose office is in the same
city as user has read only access.

if
exists(record.ManagedUsers:t1[t1.OfficeAddress.City=record.OfficeAddress.C
ity]) then

    return readOnly;
```

Note

Currently, it is not possible to:

- Use an aggregate function in between .
- Select fields of an association to aggregate values.

Operators

By default, operation evaluation order is based on precedence and associativity. The order of evaluation can be indicated explicitly by using parentheses.

The following table shows all operators from highest to lowest precedence and their associativity:

PRECEDENCE LEVEL	OPERATOR	OPERAND TYPE	RESULT TYPE	ASSOCIATIVITY
8	(access to an element of a list) . (access to fields) () (parenthesis)	List index must be a decimal.	Can be any type.	Left to right.

PRECEDENCE LEVEL	OPERATOR	OPERAND TYPE	RESULT TYPE	ASSOCIATIVITY
7	not	Boolean.	Boolean.	
6	* /	Decimal.	Decimal.	Left to right.
5	+ -	Decimal.	Decimal.	Left to right.
4	< <= > >=	String, Decimal, timestamp, date, time (3).	Boolean.	Not associative.
3	= <>	String, decimal, timestamp, date, time, boolean (3).	Boolean.	
2	and	Boolean.	Boolean.	Left to right.
1	or	Boolean.	Boolean.	Left to right.

Null value management

Arithmetic operators

An arithmetic operator (*****, **/**, **+** and **-**) returns **null** if any operand is **null** .

Comparison operators

A Comparison operator (**<** , **<=** , **>** , **>=** , **=** and **<>**) returns **null** if any operand is **null** .

Boolean operators

Boolean operators use thread-value logic.

Truth table for **and** is:

AND	TRUE	FALSE	NULL
true	true	false	null
false	false	false	false

AND	TRUE	FALSE	NULL
null	null	false	null

Truth table for **or** is:

OR	TRUE	FALSE	NULL
true	true	true	true
false	true	false	null
null	true	null	null

Index expressions

An indexed expression with an index that evaluates **null** or is out of range returns **null** .

Functions

Functions usually return **null** if a parameter is **null** . An exception is the function **isNull(value)** , which never returns **null** .

If statement

"If then" statement

An "if then" statement has the following syntax:

```
if condition-expression then
    then-body-statement
```

The condition expression must evaluate to a boolean type.

If the expression evaluates to **true** , the "then" statement is executed. If the expression evaluates to **false** or **null** , the "then" statement is ignored.

The 'then' statement is a [body statement](#) .

"If then else" statement

An "if then else" statement has the following syntax:

```
if condition-expression then  
  
    then-body-statement  
  
else  
  
    else-body-statement
```

The condition expression must be of boolean type.

If the expression evaluates to **true** , then the "then" statement is executed. If the expression evaluates to **false** or **null** , then the "else" statement is executed.

A "then" or "else" statement is a [body statement](#) .

Note

The expression:

```
if condition-expression then  
  
    statement-a;  
  
else  
  
    statement-b;
```

might not be equivalent to:

```
if not condition-expression then  
  
    statement-b;  
  
else  
  
    statement-a;
```

Indeed, if the expression is null, then the "else" statement is executed in both cases.

"Then" or "else" body statement

A body statement can be:

- A [return statement](#) ,
- An ["if then"](#) or ["if then else"](#) statement,
- A statement block.

A **statement block** is used to group multiple statements. It starts with the keyword **begin** and ends with the keyword **end** .

```
begin

    <statement 1>

    <statement 2>

    ...

    <last statement>

end
```

All statements except last one must be an ["if then"](#) or an ["if then else"](#) statement.

The last statement can be a ["if then"](#) , ["if then else"](#) or a ["return"](#) statement.

```
if isMember('sales-team')then

begin

    if record.Country='F' then

        return readWrite;

    if record.Country='UK' then

        return readOnly;

end
```

```

else

begin

    if record.Country='D' then

        return readOnly;

    if record.Country='B' then

        return readWrite;

    return hidden;

end

```

Return statement

A return statement specifies access to records that meet given conditions.

The following table shows valid return statements.

RETURN STATEMENT	DESCRIPTION
return hidden;	The record is hidden (the user has no access).
return readOnly;	The record is read only for the current user.
return readWrite;	The record can be read and modified by the current user.

If no return statement applies to a given record, the value **hidden** is assumed.

Context

Introduction

Record permissions can depend on the current dataspace, dataset, or session.

Dataspace

The predefined alias **dataspace** provides access to information on the current dataspace.

This alias gives access to the following fields:

NAME	TYPE	DESCRIPTION
name	string	The name of the dataspace. Because the dataspace namespace and the snapshot namespace are independent, the returned string is only an identifier in the context of one of the namespaces. For a global dataspace or snapshot identifier, use the field id .
id	string	The persistent identifier of a dataspace or snapshot. Compared to name , this identifier additionally specifies whether this id is for a dataspace or a snapshot.
isSnapshot	boolean	Is true if dataspace is a snapshot and false if dataspace is a branch.

Example:

```
if dataspace.name = 'branch-R' then  
  
    return readOnly;
```

Dataset

The predefined alias **dataset** provides access to information on the current dataset.

This alias has the following fields:

NAME	TYPE	DESCRIPTION
name	string	The name of the dataset.

Example:

```
if dataset.name = 'TEST' then
    return readWrite;
```

Session

The predefined alias **session** provides access to information on the current user session.

This alias has the following fields:

NAME	TYPE	DESCRIPTION
userId	string	The current user's id.
userEmail	string	The current user's email.

To check current user's roles, see [Roles](#) .

Example:

```
if session.userId = 'jdoe' then
    return readWrite;
```

Functions

Roles

The following table describes **isMember()** function:

SYNTAX	DESCRIPTION
<code>isMember('rolea', 'roleb'...)</code>	Returns true if the current user has at least one of the specified roles. The following built-in roles must be specified without quotes:

SYNTAX	DESCRIPTION
	administrator , readOnly , everyone . Custom roles are specified as string literals. Specifying a role that does not exist is not an error. (It is considered as a role with no members.) Note The roles administrator and 'administrator' can be two distinct roles. This means that isMember(administrator) might not return the same result as is_member('administrator'). The same rule applies for other built in roles.

Example:

```

if isMember('SALES', 'SUPPORT') then

    return readWrite;

// Role administrator and not surrounded by quotes:

if isMember(administrator, 'SUPPORT') then

    return readOnly;

```

String matching functions

A string matching function takes three parameters:

- A **stringExpression** that evaluates to a string.
- A **pattern** that must be a string literal.
- An optional boolean literal **isCaseSensitive** that indicates if matching is case-sensitive. If omitted, it is assumed that matching is case-insensitive.

SYNTAX	DESCRIPTION
matches(stringExpression, pattern, isCaseSensitive)	Returns true if stringExpression matches a java regular expression pattern .

SYNTAX	DESCRIPTION
<code>startsWith(stringExpression, pattern, isCaseSensitive)</code>	Returns true if stringExpression starts with string pattern .
<code>endsWith(stringExpression, pattern, isCaseSensitive)</code>	Returns true if stringExpression ends with string pattern .
<code>contains(stringExpression, pattern, isCaseSensitive)</code>	Returns true if stringExpression contains a string pattern .
<code>containsWholeWord(stringExpression, pattern, isCaseSensitive)</code>	Returns true if stringExpression contains a whole word pattern .

Examples:

```
// Give read write access if first name starts with 'a', 'b' or 'c' (case-
sensitive).

if matches(record.FirstName, '[a-c].*', true) then

    return readWrite;

// Give read only access if first name starts with 'Lé' (case-
insensitive).

if startsWith(record.FirstName, 'Lé') then

    return readOnly;

// Give read only access if first name ends with 'my' (case-insensitive).

if endsWith(record.FirstName, 'my') then

    return readOnly;

// Give read only access if email contains with 'BeauMont@' (case-
```

```

sensitive).

if contains(record.Email, 'BeauMont@', true) then

    return readOnly;

// Give read write access if last name contains the whole word
'Michel' (case-insensitive).

if containsWholeWord(record.LastName, 'Michel', false) then

    return readWrite;

```

Aggregate functions

The following aggregate function can be used with associations:

SYNTAX	DESCRIPTION
count(expression)	Returns the number of rows of an expression.
exists(expression)	Returns true if the expression evaluates to at least one row.

Examples:

```

// Give read write access if managed users count is at least 2.

if count(record.ManagedUsers[]) >= 2 then

    return readWrite;

// Give read only access if count of managed users that are not in Paris
is less than 4.

if count(record.ManagedUsers:m1[m1.OfficeAddress.City<>'Paris']) < 4 then

    return readOnly;

```

```
// Give read write access if managed users count is at least 1.

if exists(record.ManagedUsers[]) then

    return readWrite;

// Give read write access if count of managed users that are not in Briton
is at least 1.

if exists(record.ManagedUsers:u1[u1.OfficeAddress.City<>'Briton']) then

    return readWrite;
```

Miscellaneous functions

The following table describes built-in functions that return a boolean.

SYNTAX	DESCRIPTION
isNull(value)	Returns true if value is null. Value can be an expression be of any type.

Examples:

```
// Give read write access if no supervisor.

if isNull(record.supervisor) then

    return readWrite;
```

Function field

- [Introduction](#)
- [Lexical structure](#)
 - [Introduction](#)

- Character set
- Character case sensitivity
- Comments
- Keywords
- Identifiers
 - Unquoted identifier
 - Quoted identifiers
- Types
 - Simple types
 - Complex types
 - List types
 - Mutable and immutable types
 - Keyword typeof
- Literals
 - String literal
 - Decimal or integer literal
 - Timestamp literal
 - Date literal
 - Time literal
 - Boolean literal
 - Null literal
- Operators
 - Precedence
 - Arithmetic operators
 - String concatenation operator
 - Boolean operators
 - Comparison operators
 - Built-in and unit functions

- Assignments
- Statement blocks
- Variables and constants
 - Introduction
 - Assignment statements
 - Declaration with type detection
 - Explicit type definition
 - Scope
 - Complex variables
 - Multi-valued fields
- Predefined variables
 - Introduction
 - Record
 - Root
 - Dataspace
 - Dataset
- Functions and procedures
 - Functions
 - Exported function
 - Procedures
 - Parameters
- If statement
 - "If then" statements
 - "If then else" statements
- Loops
 - "For in do" loops
 - "While do" loops

- Units
- Logging and debugging.
- Initial script

Introduction

Use the function DSL (Domain Specific Language) to define a function fieldfunction field.

This DSL is a procedural and strongly statically typed language.

To create and modify a function field, use the [Data Model Assistant \(DMA\)](#) .

The Data Model Assistant (DMA) includes a script editor that provides contextual code completion.

Lexical structure

Introduction

A script has following structure:

```
<unit usage statement 1>

<unit usage statement 2>

...

<unit usage statement N>


<function or procedure definition 1>

<function or procedure definition 2>

...

<function or procedure definition M>
```

For more information, see [Unit and function and procedure](#) .

Example:


```
// This field returns the full address for current record.

export function getValue(): string

begin

    var address := record.FirstName | ' ' | record.LastName;

    for street in record.OfficeAddress.Street do

        begin

            address |= '\n' | street;

        end;

    end;

    address |= '\n' | record.OfficeAddress.ZipCode | ' ' |
record.OfficeAddress.City;

    address |= '\n' | record.OfficeAddress.Country;

    return address;

end
```

Character set

The Unicode character set is supported.

Character case sensitivity

The DSL is case-sensitive.

Comments

A single line comment extends from `//` to the end of the current line:

```
// This is a comment

if record.LastName = 'Doe' then // This is another comment.

    return true;
```

A multi-line comment extends from **/**/** and ends with ***/** :

```
/* This is an example of a multi-line

    comment */

if record.isActive then

    return false;
```

Keywords

The reserved keywords are: **and** , **or** , **not** , **uses** , **as** , **export** , **typeof** , **mutable** , **immutable** , **unmodifiable** , **function** , **procedure** , **const** , **var** , **if** , **then** , **else** , **for** , **while** , **in** , **do** , **begin** , **end** , **return** , **true** , **false** , **null** .

Reserved keywords cannot be used as plain (unquoted) identifiers.

Identifiers

Unquoted identifier

An **unquoted identifier** is an unlimited-length sequence of letters, digits, or underscore (**_**). The first character must be a letter or an underscore.

Valid letters are **a** to **z** and **A** to **Z** . Valid digits are **0** to **9** .

An unquoted identifier may not be equal to a reserved keyword.

Quoted identifiers

A quoted identifier is an unlimited length of any Unicode character except double quote (**"**).

Quoted identifiers **must** be used surrounded by double quotes.

An unquoted identifier can be used surrounded by double quotes. This means that identifier "a_name" is equal to a_name.

Quoted identifiers can be reserved keywords.

Types

Simple types

The following simple types are supported:

TYPE	KEYWORD	PROPERTIES	EBX® CORRESPONDING TYPES
Boolean	boolean		xs:boolean
Decimal (unlimited precision)	decimal		xs:decimal
Integer (32 bits)	int		xs:int xs:integer
String	string		xs:string xs:Name osd:text osd:html osd:email osd:password 1 osd:color osd:dataspaceKey osd:datasetName
Timestamp (millisecond precision and without time-zone)	timestamp	year month (1 to 12) day (1 to 31) hour (0 to 23) minute (0 to 59) second (0.000 to 59.999)	xs:dateTime
Date (without time-zone)	date	year month (1 to 12) day (1 to 31)	xs:date
Time (millisecond precision)	time	hour (0 to 23) minute (0 to 59) second (0.000 to 59.999)	xs:time

TYPE	KEYWORD	PROPERTIES	EBX® CORRESPONDING TYPES
Locale	locale		osd:locale
URI (Uniform Resource Identifier)	uri		xs:anyURI
Resource	resource		osd:resource

1 It is not possible to define a function field for type **osd:password** .

Complex types

A complex type is the type of a group (complex) node defined in an EBX® schema.

For more information, see the chapter [complex variables](#)

Reference to a schema node's type (for example, when you declare a function parameter or a variable) implies using [keyword typeof](#) .

List types

The DSL supports lists. Declare a list by using the following syntax:

```
list<item_type>
```

Declare an unmodifiable list by using following syntax:

```
unmodifiable list<item_type>
```

An item cannot be added, removed, or replaced if the list is unmodifiable. An item of the list might or might not be modifiable, depending on its type.

Use indexed notation to set or get a value. (The first index is 0):

```
// Set first value of cities.

cities[0] := 'Paris';
```

```
// Get the second item of cities.

return cities[1];
```

An index can be any expression of decimal type. Convert the index value to an integer by dropping the fractional part.

A get expression returns null if the index is null, negative, or out of range.

A set expression reports an error at runtime if the index is null, negative, or out of range.

The property **size** returns the size of the list:

```
var size := cities.size;
```

You can iterate a list. For more information, see [for loops](#).

Multi-value schema fields

Schema fields that are multi-valued are considered of list type.

A multi-valued schema field is a field that has its maximum number of values (maxOccurs) greater than one or set to "unbounded".

Indexed notation is used to access or set a value. The first index is 0. In the following example, record field 'OfficeAddress/street' is multi-valued:

```
// Return the second line of the street part of the address.

return OfficeAddress.street[1].
```

The **typeof** keyword can be used with multi-valued fields (see chapter [keyword typeof](#)).

Mutable and immutable types

You cannot modify an immutable complex or list object.

You can modify a mutable complex or list object, but other constraints may apply that can prevent modifying some attributes of an object.

A simple type is always immutable.

You can assign a mutable value to a variable or a return value of a compatible immutable type, but assigning an immutable value to a mutable variable or a return value generates an error at compile time.

Keyword `typeof`

Use the keyword **`typeof`** to specify a type depending on a field or variable. You can use this keyword preceded by the keyword **`mutable`** or **`immutable`**.

For example, to reference the type of record field **`OfficeAddress`**, use the following syntax:

```
typeof record.OfficeAddress
```

By default, the type of a record field is always immutable. You can specify a mutable type using following syntax:

```
mutable typeof record.OfficeAddress
```

If a field is multi-valued, it is considered a list. You can reference the type of an item of the list using following syntax:

```
typeof record.Addresses[*]
```

You can use the keyword 'mutable' on types that have a **`mutable`** form:

```
mutable typeof record.Addresses[*]
```

If the type does not have a mutable form, then the keyword is simply ignored. This is the case for simple types.

===

Literals

String literal

String literals can be any sequence of Unicode characters surrounded by single quotes. The following table displays characters that need to be replaced by an escape sequence:

CHARACTER	ESCAPE SEQUENCE
Tab	\t
Backspace	\b
New line	\n
Carriage return	\r
Form feed	\f
Single quote	\'
Backlash	\\

Specify a character by using a Unicode escape sequence that has format `\uXXXX` , where XXXX is the hexadecimal code of the Unicode character.

Examples

VALUE	SYNTAX
O'Harra	'O\''Harra'
Noël	'No\u00EBI'
été	'\u00e9\u00E9'

Note

An invalid escape or Unicode sequence generates an error at compile time.

Decimal or integer literal

The following decimal formats are supported:

FORMAT	EXAMPLES
Integer	546 -67
Floating point	54.987 -433.876 0.00054 -0.0032

FORMAT	EXAMPLES
Exponent notation	34.654e-5 -45E+65 1.543e23

Timestamp literal

Timestamp literals have the format **dt(yyyy-MM-dd hh:mm:ss.sss)** .

Seconds are optional. When they are not specified, 0 is assumed. Seconds can have fractions up to millisecond precision.

Any dates that are not valid in the Gregorian calendar generate errors at compile time.

Examples:

```
dt(2010-01-02 00:00:00.000)
```

```
dt(2019-2-3 12:56:7)
```

```
dt(2019-2-3 12:56:7.5)
```

```
dt(2019-5-7 1:6)
```

Date literal

Date literals have the format **d(yyyy-MM-dd)** .

Any dates that are not valid in the Gregorian calendar generate errors at compile time.

Examples:

```
d(2010-01-02)
```

```
d(2019-2-3)
```

```
dt(2019-5-7)
```

Time literal

Time literals have the format **t(hh:mm:ss.sss)** .

Seconds are optional. When they are not specified, 0 is assumed. Seconds can have fractions up to millisecond precision.

Invalid times generate an error at compile time.

Examples:

```
t(00:00:00)
```

```
t(12:56:7)
```

```
t(12:56:7.5)
```

```
t(1:6)
```

Boolean literal

A Boolean literal is either the keyword **true** or the keyword **false** .

Null literal

A null literal is the keyword **null** .

Use this literal only in the following situations:

- To set a variable or field,
- As a return value.

Note

It is not possible to use this keyword to test if a variable or a field is **null** . Instead, use the function **isNull()** .

Operators

Precedence

By default, operation evaluation order is based on precedence and associativity. Use parentheses to explicitly indicate the order of evaluation.

The following table shows all operators, from highest to lowest precedence, and their associativity:

PRECEDENCE LEVEL	OPERATOR	OPERAND TYPE	RESULT TYPE	ASSOCIATIVITY
9	(access to an element of a list) . (access to fields) () (parenthesis)	List index must be a decimal.	Can be any type.	Left to right.
8	not	Boolean	Boolean	
7	* /	Decimal	Decimal	Left to right.
6	+ -	Decimal	Decimal	Left to right.
5		(string concatenation)	String	String
4	< <= > >=	String, Decimal, timestamp, date, time (3).	Boolean	Not associative.
3	= <>	String, decimal, timestamp, date, time, boolean (3).	Boolean	
2	and	Boolean	Boolean	Left to right.
1	or	Boolean	Boolean	Left to right.

Arithmetic operators

All arithmetic operators (`*`, `/`, `+` and `-`) are evaluated in decimal with 34 digits precision (IEEE 754R Decimal128 format). Result is **null** if any operand is **null** .

A decimal value is automatically converted when assigned to an **int** type variable. An error occurs if the value cannot be converted because it is not an integer, is greater than 2147483647, or is less than -2147483648.

String concatenation operator

The string concatenation operator (`|`) operands can be of any type and are automatically converted to a string before concatenation. The **null** value is replaced by the empty string.

The string concatenation operator (`|`) replaces all null operands by the empty string before executing the concatenation.

Boolean operators

Boolean operators use thread-value logic.

The truth table for the **and** operator is as follows:

AND	TRUE	FALSE	NULL
true	true	false	null
false	false	false	false
null	null	false	null

The truth table for the **or** operator is as follows:

OR	TRUE	FALSE	NULL
true	true	true	true
false	true	false	null
null	true	null	null

Comparison operators

A comparison operator (`<` , `<=` , `>` , `=>` , `=` and `<>`) compares two operands of same type. Variables of `int` type are always converted to **decimal** before comparison.

The returned value of a comparator is of **boolean** type. This value is **true** or **false** only if all operands are not **null** . It is **null** if any of its operand is **null** .

Built-in and unit functions

Built-in and unit functions usually return **null** if a parameter is **null** .

There can be exceptions, so be sure to read the [API documentation](#) .

Assignments

The following table shows all assignment operators:

OPERATOR	OPERAND TYPE	DESCRIPTION
<code>:=</code>	Can be used with any type.	Standard assignment.
	<code>=</code>	string
<code>+=</code>	decimal	<code>a += b</code> is equivalent to <code>a := a + b</code>
<code>-=</code>	decimal	<code>a -= b</code> is equivalent to <code>a := a - b</code>
<code>*=</code>	decimal	<code>a *= b</code> is equivalent to <code>a := a * b</code>
<code>/=</code>	decimal	<code>a /= b</code> is equivalent to <code>a := a / b</code>

Statement blocks

A statement block is used to group multiple statements. It starts by the keyword **begin** and ends with the keyword **end** . Most statements need to be terminated by a `;` :

```
begin
    statement_1;
    statement_2;
end
```

A statement block can contain one or more statement blocks that can also contain statement blocks:

```
begin

    statement_1;

    begin

        statement_2_1;

        statement_2_3;

        begin

            statement_3_1;

            statement_3_2;

        end

    end

    statement_3;

end
```

Variables and constants

Introduction

Any statement in a block can declare a variable or a constant.

A variable or a constant always has a type that is fixed when the variable is declared.

A value can be assigned to a variable using the assignment operator `:=` .

Its is an error to assign a value of the wrong type to a variable, or to change the value of a constant after its declaration.

Assignment statements

Declaration with type detection

A variable or constant type can be detected automatically if initialized when declared. This feature is called type inference.

Syntax for a variable is:

```
var variable_name := a_value;
```

The syntax for a constant is:

```
const const_name := a_value;
```

Examples:

```
begin

    var firstName := 'John';           // Variable of type string.

    var age := 30;                      // Variable of type decimal.

    var address := record.address;     // Variable of an immutable complex.


    const MAX_AGE := 100;               // Decimal constant.

    const DEFAULT_COUNTRY := 'France'; // String constant.

end
```

Note

A constant is not necessarily immutable.

Explicit type definition

It is possible to explicitly specify the variable's type.

The syntax for a variable is:

```
// Following variable initial value is null.

var variable_1_name : variable_1_type;


// Value a_value must be compatible with type variable_1_type.

var variable_2_name : variable_1_type := a_value;
```

The syntax for a constant is:

```
const const_name : variable_1_type := a_value;
```

Examples:

```
begin

    var firstName : string;

    firstName := 'John';


    var age : decimal;

    age := 30;

    var address : typeof record.address;

    address := record.address;


    const MAX_AGE : decimal := 100;

    const DEFAULT_COUNTRY : string := 'France';

end
```

Variables that are not initialized have a null value:

```
begin

    var firstName : string;

    if isNull(firstName) then

        do_something(); // Statement will be executed.

    end
```

The following statements have errors:

```
begin

    // Following will not compile because variable is not initialized and no
    type

    // is defined.

    var firstName;

    // Following will not compile because variable type is decimal and value
    is

    // a string.

    var age : decimal := 'test';

end
```

Scope

A variable or constant scope spans from its declaration to the end of the statement block where it was declared:

```
begin

    var firstName := 'John';

    var lastName := 'Doe';
```



```

begin

    // Following is OK because firstName and lastName are visible.

    var fullName := firstName | ' ' | lastName;

    ....

end

// Following will NOT compile because fullName is out of the scope.

var message := 'Please contact ' | fullName;

end

```

A variable with same name cannot exist in same block or sub blocks:

```

begin

    var firstName := 'John';

    begin

        var firstName := 'Bob'; // Error! Variable already declared.

        ....

    end

end

```

The following is correct:

```

begin

  begin

    var firstName := 'Bob';

    ....

  end

  // The following is not an error, because block where previous variable
  was

  // declared is ended.

  var firstName := 'John';

end

```

Complex variables

Variable or constants of complex type can be declared using type detection or [typeof keyword](#) .

Dot notation is used to access fields of a variable of complex type.

Examples:

```

// Declaration using type detection.

var address1 := record.OfficeAddress;

// Declaration using typeof notation.

var address2 : typeof record.OfficeAddress;

address2 := address1;

```

Each step (parts separated by a dot) is an [identifier](#) . This means that following quoted notation can be used for any step:

```
var city := record."OfficeAddress".City;
```

This is useful for steps equal to a reserved keyword or using characters, such as the minus character (-) or dot (.), that are not compatible with unquoted identifiers.

At runtime, any step can evaluate to **null** . In this case the full field expression evaluates to **null** .

Multi-valued fields

Multi-valued fields are treated as ordered lists. Index notation is used to access an item of the list. Indexes are 0 based (first index is 0).

In the following example, the field "Address" is multivalued:

```
var city := record.Address[1].City;
```

If an index is out of bound, the indexed expression will return **null** . In this case the full field expression evaluates to **null** .

An index may be a decimal expression. Only the integer part of the expression will be used.

Predefined variables

Introduction

Predefined variables allow access to the context. These contextual variables are constant and of an immutable complex type.

Record

The predefined variable **record** is available only if the current script is for a function field of a table. It allows read-only access to the current record.

Its fields are defined by the current EBX® schema.

Example:

```
// Returns the full name.

export function getValue(): string

begin
```

```

        return record.FirstName | ' ' | record.LastName;

    end

```

Root

The predefined variable **root** is available only if the current script is for a function instance field of a dataset. It provides read-only access to the instance fields of current datasets.

Its fields are defined by the current EBX® schema.

Example:

```

// Returns information on current dataset data.

export function getValue(): string

begin

    return root.City | ' ' | root.Region;

end

```

Dataspace

The predefined variable **dataspace** provides access to information on the current dataspace.

This variable has following fields:

NAME	TYPE	DESCRIPTION
name	string	The name of the dataspace. Since the dataspace namespace and the snapshot namespace are independent, the returned string is only an identifier in the context of one of the namespaces. For a global dataspace or snapshot identifier, use field id .

NAME	TYPE	DESCRIPTION
id	string	The persistent identifier of a dataspace or snapshot. Compared to name , this identifier additionally specifies whether this id is for a dataspace or a snapshot.
isSnapshot	boolean	Is true if dataspace is a snapshot and false if dataspace is a branch.

Example:

```
// Returns details on the current dataspace.

export function getValue(): string
begin
    if dataspace.isSnapshot then
        return 'Snapshot: ' | dataspace.name;
    else
        return 'Branch: ' | dataspace.name;
    end
end
```

Dataset

The predefined variable **dataset** provides access to information on the current dataset.

This variable has the following fields:

NAME	TYPE	DESCRIPTION
name	string	The name of the dataset

Example:

```
// Returns the current dataset name.

export function getValue(): string

begin

    return dataset.name;

end
```

Functions and procedures

Functions

Functions are methods that return a value.

A function with no parameters has the following syntax:

```
function function_name(): return_value_type

begin

    ....

    return a_value;

end
```

A function with parameters has the following syntax:

```
function function_name(

    parameter_1 : parameter_1_type,

    parameter_2 : parameter_2_type...): return_value_type

begin

    ....

    return a_value;
```

```
end
```

The last statement of a function must always be a return statement. The function can include as many return statement as necessary.

Example:

```
function getFullName(firstName: string, lastName: string): string
begin
    if isNull(firstName) then
        return lastName;

    return firstName | ' ' | lastName;
end
```

It is illegal to return a value different from the one declared:

```
function getValue(): string
begin
    return 0; // Error! Return type is decimal, not string.
end
```

Exported function

An exported function can be called directly by EBX®.

Only one exported function is allowed per script. Its signature (name, and return type) depends on the type of the function field.

For the field of a record, definition must be similar to:

```

export function getValue(): typeof record.<path>

begin

    ...

end

```

For the field of a dataset, definition must be similar to:

```

export function getValue(): typeof root.<path>

begin

    ...

end

```

When a function field is first edited, EBX® provides an [initial script](#) with the correct definition of the exported function.

Procedures

Procedures are methods that do not return a value.

A procedure with no parameters has the following syntax:

```

procedure procedure_name()

begin

    ....

end

```

A procedure with parameters has the following syntax:

```

procedure procedure_name(

    parameter_1 : parameter_1_type,

```



```

    parameter_2 : parameter_2_type...)

begin

    ....

end

```

A procedure cannot have a return statement, but the function can include as many return statements as necessary. A procedure statement cannot return a value.

Parameters

Simple type parameters are passed by value. This means that a function or a procedure receives a copy of the original value.

Complex and list types are passed by reference. This means that a function or a procedure receives the original object. If the function or procedure modifies the object, the original one is modified.

If statement

"If then" statements

An "if then" statement has the following syntax:

```

if condition-expression then

    then-statement

```

The condition expression must evaluate to a boolean type.

The 'then' statement can be a [statement block](#) .

"If then else" statements

An "if then else" statement has following syntax:

```

if condition-expression then

    then-statement

```

```

else

    else-statement

```

The condition expression must be of boolean type.

A 'then' or 'else' statement can be a [statement block](#) .

Note

Expression:

```

if condition-expression then

    statements-a;

else

    statements-b;

```

Cannot be equivalent to:

```

if not condition-expression then

    statements-b;

else

    statements-a;

```

Indeed, if the expression is null, the **else** statement is executed in both cases.

Loops

"For in do" loops

A "for in do" loop statement is used to select each item of a list and execute a block statement. It has following syntax:

```

for item_variable_name in list do

begin

    statement_a;

    statement_b;

    ...

end

```

The block statement is executed once for each value of the list. At each iteration, the read-only item variable takes a value of the list in the order of the list.

The name of the item variable must be unique in the current scope or an error will be generated at compile time.

If a single statement must be executed for each item, the following simpler syntax can be used:

```

for item_variable_name in list do statement;

```

The following example iterates a list of complex:

```

// Concatenate all city addresses in a single string.

var cities := '';

for address in record.Addresses do

begin

    if cities <> '' then cities |= ', ';

    cities |= address.city;

end

```

"While do" loops

A "while do" loop statement is used to execute a statement block until a condition is **true** . It has following syntax:

```
while condition do  
  
begin  
  
    statement_a;  
  
    statement_b;  
  
    ...  
  
end
```

If a single statement must be executed, the following simpler syntax can be used:

```
while condition do statement_a;
```

The following example calculates factorial of a value:

```
function factorial(value : decimal): decimal  
  
begin  
  
    var factorial := value;  
  
    while(value > 1) do  
  
        begin  
  
            value -= 1;  
  
            factorial *= value;  
  
        end  
  
    return factorial;
```

```
end
```

Units

EBX® provides an API that is packaged in "units".

A unit can define multiple function or procedures.

Except for the default one, a unit must be declared before usage. The declaration must be at the top of the script and must follow these types of syntaxes:

```
uses package_name.unit_name_a;

uses package_name.unit_name_b as alias_b;
```

Currently, **package_name** is always **core** . A unit alias must be unique in a script.

Methods can be referenced using following syntaxes:

```
value1 := package_name.unit_name_a.function_a();

value2 := package_name.unit_name_a.function_b(parameter1, parameter2);

value3 := alias_a.function_c();

value4 := alias_a.function_d(parameter1, parameter2);

package_name.unit_name_a.procedure_a();

package_name.unit_name_a.procedure_b(parameter1, parameter2);

alias_a.procedure_c();

alias_a.procedure_d(parameter1, parameter2);
```

The following example uses the **core.list** unit to create a list:

```
uses core.list as list;

export function getValue(): typeof record.Cities

begin

    return list.of('Paris', 'Bruxelles', 'Berlin');

end
```

For details on the provided units, see the [API Documentation](#) .

Logging and debugging.

The unit [unit.log](#) provides a function that can be used to log a message.

A message logged by scripts can be viewed using **Administration>Repository management>Scripting** EBX® menu.

Another useful feature is that if a runtime encounters an error while executing a script, it is usually logged with the line in the script where the error occurred.

Initial script

When a new function field is created using the [DMA](#) , an initial script is created.

The following example is a script created for a field of type string:

```
export function getValue(): string

begin

    return 'A string value';

end
```

The exported function signature (name, and return type) depends on the field's type and should not be changed.

Function field API

- [Unit summary](#)
- [Unit default](#)
- [Unit core.complex](#)
- [Unit core.date](#)
- [Unit core.datetime](#)
- [Unit core.list](#)
- [Unit core.locale](#)
- [Unit core.log](#)
- [Unit core.math](#)
- [Unit core.resource](#)
- [Unit core.string](#)
- [Unit core.time](#)
- [Unit core.uri](#)

Unit summary

UNITS	DESCRIPTION
default	The default unit.
core.complex	Script unit that provides methods to create and update complex values.
core.date	Script unit that provides date functions.
core.datetime	Script unit that provides datetime functions.
core.list	Script unit that provides methods to create and update lists.
core.locale	Script unit that provides functions for managing locales.
core.log	Script unit that provides methods to log messages.

UNITS	DESCRIPTION
core.math	Script unit that provides mathematical functions and constants.
core.resource	This script unit methods that generates resource references compatible with the schema type <code>ods:resource</code> .
core.string	Script unit that provides string manipulation functions.
core.time	Script unit that provides time functions.
core.uri	Script unit that provides functions for managing Uniform Resource Identifier (URI) references.

Unit default

The default unit. This unit is automatically included in a script.

Methods from this script can be called directly and do not require a *uses* statement.

METHODS
function <code>isNull(value: any_type): boolean</code> Returns <code>true</code> if value is null.
function <code>isTrueOrNull(value: boolean): boolean</code> Returns <code>true</code> if value is true or null.
function <code>isFalseOrNull(value: boolean): boolean</code> Returns <code>true</code> if value is false of null.

- function `isNull(value: any\_type): boolean`
- function `isTrueOrNull(value: boolean): boolean`
- function `isFalseOrNull(value: boolean): boolean`

function `isNull(value: any\_type): boolean`

Returns `true` if value is null.

Parameters :

value a value of any type.

Return :

`true` if `value` is null or else `false` . Never returns null.

function isTrueOrNull(value: boolean): boolean

Returns `true` if value is true or null.

Parameters :

value: a boolean value.

Return :

`true` if `value` is null or `true` , or else `false` . Never returns null.

function isFalseOrNull(value: boolean): boolean

Returns `true` if value is false of null.

Parameters :

value: a boolean value.

Return :

`true` if `value` is null or `false` , or else `false` . Never returns null.

Unit core.complex

Script unit that provides methods to create and update complex values.

METHODS
<code>of<complexType>(): complexType</code> Create an instance of a complex type.
<code>of<foreignKeyType>(): foreignKeyType</code> Create an instance of a foreign key type.

- `of<complexType>(): complexType`
- `of<foreignKeyType>(): foreignKeyType`

of<complexType>(): complexType

Create an instance of a complex type.

Example:

```

uses core.complex as complex;

export function getValue(): typeof record.OfficeAddress
begin
    var value := complex.of<typeof record.OfficeAddress>();

    value.Street := '4323 Broadway';

    value.City := 'New York';

    value.State := 'NY';

    value.Zip := '10019';

    value.Country := 'USA';

    return value;
end

```

Function Types :

complexType: the object type. Is mandatory. Should be an expression `typeof identifier`.

Return :

the new object. Is always the mutable version of the specified type.

of<foreignKeyType>(): foreignKeyType

Create an instance of a foreign key type.

Example:

```

uses core.complex as complex;

export function getValue(): typeof record.Supervisor

```

```

begin

    var foreignKey := complex.foreignKeyOf<typeof record.Supervisor>();

    foreignKey.Id := 435;

    return foreignKey;

end

```

Function Types :

foreignKeyType: the foreign key type. Is mandatory. Should be an expression `typeof identifier`.

Return :

the new foreign key object. Is always mutable.

Unit core.date

Script unit that provides date functions.

METHODS
function addYears(value: date, years: decimal): date Adds years to a date.
function addMonths(value: date, months: decimal): datetime Adds months to a date.
function addDays(value: date, days: decimal): date Adds days to a date.
function fromDatetime(value: datetime): date Converts a date and time to a date.
function days(startValue: date, endValue: date): decimal Returns the number of days between two dates.
function seconds(startValue: date, endValue: date): decimal Returns the number of seconds between two dates.

- function addYears(value: date, years: decimal): date
- function addMonths(value: date, months: decimal): datetime
- function addDays(value: date, days: decimal): date

- function fromDatetime(value: datetime): date
- function days(startValue: date, endValue: date): decimal
- function seconds(startValue: date, endValue: date): decimal

function addYears(value: date, years: decimal): date

Adds years to a date.

Parameters :

value: the input date

years: the number of years to be added. Must be an integer.

Return :

the new date or null if any parameter is null.

function addMonths(value: date, months: decimal): datetime

Adds months to a date.

Parameters :

value: the input date

months: the number of months to be added. Must be an integer.

Return :

the new date or null if any parameter is null.

function addDays(value: date, days: decimal): date

Adds days to a date.

Parameters :

value: the input date

months: the number of days to be added. Must be an integer.

Return :

the new date or null if any parameter is null.

function fromDatetime(value: datetime): date

Converts a date and time to a date. The input time part is ignored.

Parameters :

value: the input date and time

Return :

the new date or null if parameter is null.

function days(startValue: date, endValue: date): decimal

Returns the number of days between two dates.

Parameters :

startValue: the input start date

endValue: the input end date

Return :

the number of days or null if any parameter is null.

function seconds(startValue: date, endValue: date): decimal

Returns the number of seconds between two dates.

Parameters :

startValue: the input start date inclusive

endValue: the input end date exclusive

Return :

the number of seconds or null if any parameter is null.

Unit core.datetime

Script unit that provides datetime functions.

METHODS

function addYears(value: datetime, years: decimal): datetime Adds years to a date and time.

METHODS

function addMonths(value: datetime, months: decimal): datetime Adds months to a date and time.

function addDays(value: datetime, days: decimal): datetime Adds days to a date and time.

function addHours(value: datetime, hours: decimal): datetime Adds hours to a date and time.

function addMinutes(value: datetime, minutes: decimal): datetime Adds minutes to a date and time.

function addSeconds(value: datetime, seconds: decimal): datetime Adds seconds to a date and time.

function of(year: decimal, month: decimal, day: decimal, hour: decimal, minute: decimal, second: decimal): datetime Creates a date and time given year, month, day, hours, minutes and seconds.

function fromDate(value: date): datetime Converts a date to date and time with time set to 00:00:00.

function fromDateAndTime(dateValue: date, timeValue: time): datetime Creates a date and time given date and time values.

function days(startValue: datetime, endValue: datetime): decimal Returns the number of days between two date and times.

function seconds(startValue: datetime, endValue: datetime): decimal Returns the number of seconds between two date and times.

- function addYears(value: datetime, years: decimal): datetime
- function addMonths(value: datetime, months: decimal): datetime
- function addDays(value: datetime, days: decimal): datetime
- function addHours(value: datetime, hours: decimal): datetime
- function addMinutes(value: datetime, minutes: decimal): datetime
- function addSeconds(value: datetime, seconds: decimal): datetime
- function of(year: decimal, month: decimal, day: decimal, hour: decimal, minute: decimal, second: decimal): datetime
- function fromDate(value: date): datetime
- function fromDateAndTime(dateValue: date, timeValue: time): datetime
- function days(startValue: datetime, endValue: datetime): decimal

- function seconds(startValue: datetime, endValue: datetime): decimal

function addYears(value: datetime, years: decimal): datetime

Adds years to a date and time.

Parameters :

value: the input date and time

months: the number of years to be added. Must be an integer.

Return :

the new date and time or null if any parameter is null.

function addMonths(value: datetime, months: decimal): datetime

Adds months to a date and time.

Parameters :

value: the input date and time

months: the number of months to be added. Must be an integer.

Return :

the new date and time or null if any parameter is null.

function addDays(value: datetime, days: decimal): datetime

Adds days to a date and time.

Parameters :

value: the input date and time

months: the number of days to be added. Must be an integer.

Return :

the new date and time or null if any parameter is null.

function addHours(value: datetime, hours: decimal): datetime

Adds hours to a date and time.

Parameters :

value: the input date and time

hours: the number of hours to be added. Must be an integer.

Return :

the new date and time or null if any parameter is null.

function addMinutes(value: datetime, minutes: decimal): datetime

Adds minutes to a date and time.

Parameters :

value: the input date and time

minutes: the number of minutes to be added. Must be an integer.

Return :

the new date and time or null if any parameter is null.

function addSeconds(value: datetime, seconds: decimal): datetime

Adds seconds to a date and time. Fractions of seconds are supported with millisecond precision.

Parameters :

value: the input date and time

seconds: the number of seconds to be added. Fractions are rounded to three decimal digits.

Return :

the new date and time or null if any parameter is null.

function of(year: decimal, month: decimal, day: decimal, hour: decimal, minute: decimal, second: decimal): datetime

Creates a date and time given year, month, day, hours, minutes and seconds. Fractions of seconds are supported with millisecond precision.

Parameters :

year: the input year

month: the input month

day: the input day

hour: the input hour

minute: the input minute

second: the input second. Fractions are rounded to three decimal digits.

Return :

the new date and time or null if any parameter is null.

function fromDate(value: date): datetime

Converts a date to date and time with time set to 00:00:00.

Parameters :

value: the input date

Return :

the new date and time or null if parameter is null.

function fromDateAndTime(dateValue: date, timeValue: time): datetime

Creates a date and time given date and time values.

Parameters :

dateValue: the input date

timeValue: the input time

Return :

the new date and time or null if any parameter is null.

function days(startValue: datetime, endValue: datetime): decimal

Returns the number of days between two date and times. This function is equivalent to equivalent to days(toDate(startValue), toDate(endValue)).

Parameters :

startValue: the input start date and time

endValue: the input end date and time

Return :

the number of days or null if any parameter is null.

function seconds(startValue: datetime, endValue: datetime): decimal

Returns the number of seconds between two date and times. Fractions of seconds are supported with millisecond precision.

Parameters :

startValue: the input start date and time inclusive

endValue: the input end date and time exclusive

Return :

the number of seconds or null if any parameter is null.

Unit core.list

Script unit that provides methods to create and update lists.

METHODS
function of<ItemType>(value: ItemType [, value: ItemType]...): list<ItemType> Creates a list with the elements passed as parameters.
function isEmpty(value: list<ItemType>): boolean Returns <code>true</code> if value is null or a zero size list.
function contains(values: list<ItemType>, element: ItemType): boolean Returns <code>true</code> if <code>values</code> list contains an <code>element</code> .
function containsAll(values: list<ItemType>, element: ItemType): boolean Returns <code>true</code> if <code>values</code> list contains all <code>elements</code> .
procedure add(values: list<ItemType>, element: ItemType) Add an <code>element</code> to <code>values</code> list.
function remove(values: list<ItemType>, element: ItemType): boolean Remove the first occurrence of an <code>element</code> from <code>values</code> list.
procedure addAll(values: list<ItemType>, elements: list<ItemType>): boolean Add all <code>elements</code> to <code>values</code> list.

METHODS

function removeAll(values: list<ItemType>, elements: list<ItemType>): boolean Remove all `elements` from `values` list.

function retainAll(values: list<ItemType>, elements: list<ItemType>): boolean Retains only the elements in the `values` list that are contained in the `elements`.

function indexOf(values: list<ItemType>, element: ItemType) Return the index of the first occurrence of `element` in `values` list.

function lastIndexOf(values: list<ItemType>, element: ItemType) Return the index of the last occurrence of `element` in `values` list.

function subList<ItemType>(values: list<ItemType>, startIndex: decimal, endIndex: decimal): list<ItemType> Returns a new list that contains a subsequence of elements from `values` list.

procedure sort(values: list<ItemType>) Sorts the specified `values` list into ascending order, according to the natural ordering of its elements.

procedure reverseSort(values: list<ItemType>) Sorts the specified `values` list into descending order, according to the natural ordering of its elements.

procedure reverse(values: list<ItemType>) Reverses the order of elements in `values`.

- function of<ItemType>(value: ItemType [, value: ItemType]...): list<ItemType>
- function isEmpty(value: list<ItemType>): boolean
- function contains(values: list<ItemType>, element: ItemType): boolean
- function containsAll(values: list<ItemType>, element: ItemType): boolean
- procedure add(values: list<ItemType>, element: ItemType)
- function remove(values: list<ItemType>, element: ItemType): boolean
- procedure addAll(values: list<ItemType>, elements: list<ItemType>): boolean
- function removeAll(values: list<ItemType>, elements: list<ItemType>): boolean
- function retainAll(values: list<ItemType>, elements: list<ItemType>): boolean
- function indexOf(values: list<ItemType>, element: ItemType)
- function lastIndexOf(values: list<ItemType>, element: ItemType)
- function subList<ItemType>(values: list<ItemType>, startIndex: decimal, endIndex: decimal): list<ItemType>

- procedure sort(values: list<ItemType>)
- procedure reverseSort(values: list<ItemType>)
- procedure reverse(values: list<ItemType>)

function of<ItemType>(value: ItemType ...): list<ItemType>

Creates a list with the elements passed as parameters. All elements must be of same type.

Example:

```
uses core.list as list;

function getCities(): list<string>
begin
    return list.of('Paris', 'Bruxelles', 'Berlin');
end
```

Function Types :

ItemType: the type of an item of the returned list. Is optional if at least one value is specified.

Parameters :

value: an item to add to the list. All items must be of same type.

Return :

the new list. It is updatable.

function isEmpty(value: list<ItemType>): boolean

Returns `true` if value is null or a zero size list.

Parameters :

values: an input list.

Return :

`true` if `value` is null or a zero size list, or else `false`. Never returns null.

function contains(values: list<ItemType>, element: ItemType): boolean

Returns `true` if `values` list contains an `element`.

Parameters :

values: the input list. Contents must be of the same kind then `element`.

element: the input element to look for. Must be of the same kind than `values` > contents.

Return :

`true` if `value` contains the `element` else `false`.

Returns `null` if values is null.

function containsAll(values: list<ItemType>, element: ItemType): boolean

Returns `true` if `values` list contains all `elements`.

Parameters :

values: the input list. Contents must be of the same kind then `element`.

elements: the input elements to look for. Must be of the same kind than `values` >.

Return :

`true` if `value` contains all `elements` else `false`.

Returns `null` if any parameters is null.

procedure add(values: list<ItemType>, element: ItemType)

Add an `element` to `values` list.

Parameters :

values: the input list.

element: the input element to add. Its type must be compatible with the type of the list.

function remove(values: list<ItemType>, element: ItemType): boolean

Remove the first occurrence of an `element` from `values` list.

Parameters :

values: the input list.

element: the input element to add. Its type must be compatible with the type of the list.

Return :

`true` if `values` list contents was modified by this call else `false`.

procedure addAll(values: list<ItemType>, elements: list<ItemType>): boolean

Add all `elements` to `values` list.

Parameters :

values: the input list.

elements: the elements to add. The element type must be compatible with the type of the list.

Return :

`true` after `values` list contents was modified by this call else `false`.

function removeAll(values: list<ItemType>, elements: list<ItemType>): boolean

Remove all `elements` from `values` list.

Parameters :

values: the input list.

elements: the elements to remove. The element type must be compatible with the type of the list.

Return :

`true` if `values` list contents was modified by this call else `false`.

function retainAll(values: list<ItemType>, elements: list<ItemType>): boolean

Retains only the elements in the `values` list that are contained in the `elements`. In other words, removes from the `values` list all of its elements that are not contained in the `elements`.

Parameters :

values: the input list.

elements: the elements to retain. The element type must be compatible with the type of the list.

Return :

`true` if `values` list contents was modified by this call else `false` .

function indexOf(values: list<ItemType>, element: ItemType)

Return the index of the first occurrence of `element` in `values` list.

Parameters :

values: the input list.

element: the input element to look for. Its type must be compatible with the type of the list.

Return :

the index of the first occurrence of `element` in `values` list or `null` if not found.

function lastIndexOf(values: list<ItemType>, element: ItemType)

Return the index of the last occurrence of `element` in `values` list.

Parameters :

values: the input list.

element: the input element to look for.

Return :

the index of the last occurrence of `element` in `values` list or `null` if not found.

function subList<ItemType>(values: list<ItemType>, startIndex: decimal, endIndex: decimal): list<ItemType>

Returns a new list that contains a subsequence of elements from `values` list.

The subList start from the specified startIndex and extends to the endIndex of the input list.

Function Types :

ItemType: the type of an item of the returned list. It is optional.

Parameters :

values: the input list

startIndex: the input start index. Must be a positive integer or zero.

endIndex: the input end index. Must be a positive integer or zero.

Return :

the new list or null if any parameter is null.

procedure sort(values: list<ItemType>)

Sorts the specified `values` list into ascending order, according to the natural ordering of its elements. Does nothing if `values` is null or a zero size list.

Parameters :

values: the input list to sort

procedure reverseSort(values: list<ItemType>)

Sorts the specified `values` list into descending order, according to the natural ordering of its elements. Does nothing if `values` is null or a zero size list.

Parameters :

values: the input list to sort.

procedure reverse(values: list<ItemType>)

Reverses the order of elements in `values`. Does nothing if `values` is null or a zero size list.

Parameters :

values: the input list to reverse.

Unit core.locale

Script unit that provides functions for managing locales.

METHODS
<code>function of(language: string, country: string): locale</code> Creates a locale.
<code>function fromString(value: string): locale</code> Creates a locale from a string.

- `function of(language: string, country: string): locale`
- `function fromString(value: string): locale`

function of(language: string, country: string): locale

Creates a locale.

Parameters :

language: An ISO 639 alpha-2 or alpha-3 language code.

country: An ISO 3166 alpha-2 country code or a UN M.49 numeric-3 area code.

Return :

the locale.

function fromString(value: string): locale

Creates a locale from a string.

Parameters :

value: A locale, for example 'en-US'.

Return :

the locale.

Unit core.log

Script unit that provides methods to log messages.\n\nThis message can be viewed using administration user interface provide by EBX.

METHODS
<code>procedure error(message: string)</code> Logs an error message.
<code>procedure warn(message: string)</code> Logs a warning message.
<code>procedure info(message: string)</code> Logs an information message.

- `procedure error(message: string)`
- `procedure warn(message: string)`

- `procedure info(message: string)`

procedure error(message: string)

Logs an error message.

Parameters :

value: message: the message.

procedure warn(message: string)

Logs a warning message.

Parameters :

value: message: the message.

procedure info(message: string)

Logs an information message.

Parameters :

value: message: the message.

Unit core.math

Script unit that provides mathematical functions and constants.

METHODS
<code>function sin(value: decimal): decimal</code> Returns the trigonometric sine of an angle.
<code>function cos(value: decimal): decimal</code> Returns the trigonometric cosine of an angle.
<code>function tan(value: decimal): decimal</code> Returns the trigonometric tangent of an angle.
<code>function asin(value: decimal): decimal</code> Returns the arc sine of a value.
<code>function acos(value: decimal): decimal</code> Returns the arc cosine of a value.
<code>function atan(value: decimal): decimal</code> Returns the arc tangent of a value.
<code>function toDegrees(value: decimal): decimal</code> Converts an angle measured in degrees to an approximately equivalent angle measured in radians.

METHODS

function toRadians(value: decimal): decimal Converts an angle measured in radians to an approximately equivalent angle measured in degrees.

function exp(value: decimal): decimal Returns Euler's number e raised to the power of a decimal value.

function log(value: decimal): decimal Returns the natural logarithm (base e) of the input value.

function scaleByPowerOfTen(value: decimal, power: decimal): decimal Returns the value scale by power of ten.

function sqrt(value: decimal): decimal Returns the positive square root of a value.

function cbrt(value: decimal): decimal Returns the cube root of a decimal value.

function abs(value: decimal): decimal Returns the absolute value of a decimal value.

function round(value: decimal): decimal Returns the integer round value of the input value.

function roundHalfEven(value: decimal): decimal Returns the integer round half even value of the input value.

function roundHalfDown(value: decimal): decimal Returns the integer round half down value of the input value.

function roundUp(value: decimal): decimal Returns the integer round up value of the input value.

function roundDown(value: decimal): decimal Returns the integer round down value of the input value.

function floor(value: decimal): decimal Returns the integer floor value of the input value.

function ceil(value: decimal): decimal Returns the integer ceil value of the input value.

function pi(): decimal Returns pi value: 3.

function e(): decimal Returns e value: 2.

- function sin(value: decimal): decimal
- function cos(value: decimal): decimal
- function tan(value: decimal): decimal

- function `asin(value: decimal): decimal`
- function `acos(value: decimal): decimal`
- function `atan(value: decimal): decimal`
- function `toDegrees(value: decimal): decimal`
- function `toRadians(value: decimal): decimal`
- function `exp(value: decimal): decimal`
- function `log(value: decimal): decimal`
- function `scaleByPowerOfTen(value: decimal, power: decimal): decimal`
- function `sqrt(value: decimal): decimal`
- function `cbrt(value: decimal): decimal`
- function `abs(value: decimal): decimal`
- function `round(value: decimal): decimal`
- function `roundHalfEven(value: decimal): decimal`
- function `roundHalfDown(value: decimal): decimal`
- function `roundUp(value: decimal): decimal`
- function `roundDown(value: decimal): decimal`
- function `floor(value: decimal): decimal`
- function `ceil(value: decimal): decimal`
- function `pi(): decimal`
- function `e(): decimal`

function `sin(value: decimal): decimal`

Returns the trigonometric sine of an angle.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Special cases:

If the input value is zero, then the result is a zero.

Parameters :

value: the input value in radians.

Return :

the sine of the input value or null if the input value is null.

function cos(value: decimal): decimal

Returns the trigonometric cosine of an angle.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value in radians.

Return :

the cosine of the input value or null if the input value is null.

function tan(value: decimal): decimal

Returns the trigonometric tangent of an angle.

Special cases:

If the input value is zero, then the result is a zero.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value in radians.

Return :

the tangent of the input value or null if the input value is null.

function asin(value: decimal): decimal

Returns the arc sine of a value.

The returned angle is in the range $-pi/2$ through $pi/2$.

Special cases:

If the input value is zero, then the result is a zero.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value. Absolute value must not be greater than 1.

Return :

the arc sine of the input value or null if the input value is null.

function acos(value: decimal): decimal

Returns the arc cosine of a value.

The returned angle is in the range 0.0 through π .

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value. Absolute value must not be greater than 1.

Return :

the arc cosine of the input value or null if the input value is null.

function atan(value: decimal): decimal

Returns the arc tangent of a value.

The returned angle is in the range $-\pi/2$ through $\pi/2$.

Special cases:

If the input value is zero, then the result is a zero.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value.

Return :

the arc tangent of the input value or null if the input value is null.

function toDegrees(value: decimal): decimal

Converts an angle measured in degrees to an approximately equivalent angle measured in radians.

The conversion from degrees to radians is generally inexact.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value in radians

Return :

the measurement of the input value in degrees or null if the input value is null.

function toRadians(value: decimal): decimal

Converts an angle measured in radians to an approximately equivalent angle measured in degrees.

The conversion from radians to degrees is generally inexact;

users should *not* expect `cos(toRadians(90.0))` to exactly equal `0.0`.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value in degrees

Return :

the measurement of the input value in radians or null if the input value is null.

function exp(value: decimal): decimal

Returns Euler's number e raised to the power of a `decimal` value.

Calculations are executed after conversion in double and use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value.

Return :

the value e^{value} where e is the base of the natural logarithms.

If the input value is null returns null.

function log(value: decimal): decimal

Returns the natural logarithm (base e) of the input value.

Calculations use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value. Must be strictly greater than zero.

Return :

the natural logarithm of the input value or null if input value is null.

function scaleByPowerOfTen(value: decimal, power: decimal): decimal

Returns the `value` scale by `power` of ten.

Parameters :

value: the input value.

power: the power of ten that the value will be scale to to. Must be an integer.

Return :

the input value scaled by the `power` of ten or null if input value or power is null.

function sqrt(value: decimal): decimal

Returns the positive square root of a value.

Calculations use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value. Must be positive or null.

Return :

the positive square root of the input value or null if input value is null.

function cbrt(value: decimal): decimal

Returns the cube root of a `decimal` value. Calculations use double-precision 64-bit IEEE 754 floating point.

Parameters :

value: the input value.

Return :

the cube root of the input value or null if input value is null.

function abs(value: decimal): decimal

Returns the absolute value of a `decimal` value.

Parameters :

value: the input value.

Return :

the absolute value of the input value or null if input value is null.

function round(value: decimal): decimal

Returns the integer round value of the input value. This method uses the standard round half up rounding mode.

Parameters :

value: the input value.

Return :

the round half up value of the input value or null if input value is null.

function roundHalfEven(value: decimal): decimal

Returns the integer round half even value of the input value.

Parameters :

value: the input value.

Return :

the round half even value of the input value or null if input value is null.

function roundHalfDown(value: decimal): decimal

Returns the integer round half down value of the input value.

Parameters :

value: the input value.

Return :

the round half down value of the input value or null if input value is null.

function roundUp(value: decimal): decimal

Returns the integer round up value of the input value.

Parameters :

value: the input value.

Return :

the round up value of the input value or null if input value is null.

function roundDown(value: decimal): decimal

Returns the integer round down value of the input value.

Parameters :

value: the input value.

Return :

the round down value of the input value or null if input value is null.

function floor(value: decimal): decimal

Returns the integer floor value of the input value.

Parameters :

value: the input value.

Return :

the floor value of the input value or null if input value is null.

function ceil(value: decimal): decimal

Returns the integer ceil value of the input value.

Parameters :

value: the input value.

Return :

the ceil value of the input value or null if input value is null..

function pi(): decimal

Returns `pi` value: 3.141592653589793238462643383279503

Return :

pi value.

function e(): decimal

Returns `e` value: 2.718281828459045235360287471352662

Return :

e value.

Unit core.resource

This script unit methods that generates resource references compatible with the schema type

`ods:resource` .

Example:

```
uses core.resource as resource;

export function getValue(): typeof record.functions.ResourceValue
```

```

begin

    return resource.toImage('ebx-test', 'on_anim_wait.gif');

end

```

METHODS

function toImage(moduleName: string, filePath: string): string Returns a reference to an image resource.

function toIcon(moduleName: string, filePath: string): string Returns a reference to an icon resource.

function toJavaScript(moduleName: string, filePath: string): string Returns a reference to an JavaScript file resource.

function toStyleSheet(moduleName: string, filePath: string): string Returns a reference to an style sheet file resource.

function toHtml(moduleName: string, filePath: string): string Returns a reference to an image resource.

- **function toImage(moduleName: string, filePath: string): string**
- **function toIcon(moduleName: string, filePath: string): string**
- **function toJavaScript(moduleName: string, filePath: string): string**
- **function toStyleSheet(moduleName: string, filePath: string): string**
- **function toHtml(moduleName: string, filePath: string): string**

function toImage(moduleName: string, filePath: string): string

Returns a reference to an image resource.

Parameters :

moduleName: the module name.

filePath: the local file path for the resource.

Return :

the resource string compatible with the schema type `ods:resource`.

function tolcon(moduleName: string, filePath: string): string

Returns a reference to an icon resource.

Parameters :

moduleName: the module name.

filePath: the local file path for the resource.

Return :

the resource string compatible with the schema type `ods:resource`.

function toJavaScript(moduleName: string, filePath: string): string

Returns a reference to an JavaScript file resource.

Parameters :

moduleName: the module name.

filePath: the local file path for the resource.

Return :

the resource string compatible with the schema type `ods:resource`.

function toStyleSheet(moduleName: string, filePath: string): string

Returns a reference to an style sheet file resource.

Parameters :

moduleName: the module name.

filePath: the local file path for the resource.

Return :

the resource string compatible with the schema type `ods:resource`.

function toHtml(moduleName: string, filePath: string): string

Returns a reference to an image resource.

Parameters :

moduleName: the module name.

filePath: the local file path for the resource.

Return :

the resource string compatible with the schema type `ods:resource` .

Unit core.string

Script unit that provides string manipulation functions.

METHODS
function isEmpty(value: string): boolean Returns <code>true</code> if value is null or a zero length string.
function right(value: string, index: decimal): string Returns a new string that contains the right characters from <code>value</code> string.
function substring(value: string, startIndex: decimal, endIndex: decimal): string Returns a new string that contains a subsequence of characters from <code>value</code> string.
function toUpperCase(value: string): string Returns a new string with all characters from input <code>value</code> string to upper case.
function toLowerCase(value: string): string Returns a new string with all characters from input <code>value</code> string to lower case.
function replaceAll(value: string, pattern: string, replacement: string): string Returns a new string with all <code>pattern</code> occurrences in input <code>value</code> string replaced with <code>replacement</code> .
function replaceFirst(value: string, pattern: string, replacement: string): string Returns a new string with the first pattern occurrence in input <code>value</code> string replaced with replacement.
function startsWith(value: string, prefix: string): boolean Returns <code>true</code> if a <code>value</code> string starts with a prefix.
function endsWith(value: string, suffix: string): boolean Returns <code>true</code> if a <code>value</code> string ends with a suffix.
function contains(value: string, searchString: string): boolean Returns <code>true</code> if a <code>value</code> string contains a <code>searchString</code> .

METHODS

function join(separator: string [, substring: string]...): string Returns a new string composed of all input substrings separated by as separator string.

- **function isEmpty(value: string): boolean**
- **function right(value: string, index: decimal): string**
- **function substring(value: string, startIndex: decimal, endIndex: decimal): string**
- **function toUpperCase(value: string): string**
- **function toLowerCase(value: string): string**
- **function replaceAll(value: string, pattern: string, replacement: string): string**
- **function replaceFirst(value: string, pattern: string, replacement: string): string**
- **function startsWith(value: string, prefix: string): boolean**
- **function endsWith(value: string, suffix: string): boolean**
- **function contains(value: string, searchString: string): boolean**
- **function join(separator: string [, substring: string]...): string**

function isEmpty(value: string): boolean

Returns `true` if value is null or a zero length string.

Parameters :

value: a string value.

Return :

`true` if `value` is null or a zero length string, or else `false`. Never returns null.

function right(value: string, index: decimal): string

Returns a new string that contains the right characters from `value` string.

The starting starts from the specified index and extends to the end of the string.

Parameters :

value: the input string

index: the input start index. Must be a integer greater or equal to zero and less than the length of the string.

Return :

the new string or null if any parameter is null.

function substring(value: string, startIndex: decimal, endIndex: decimal): string

Returns a new string that contains a subsequence of characters from `value` string.

The sub string start from the specified startIndex and extends to the endIndex of the input string.

Parameters :

value: the input string

startIndex: the input start index. Must be a integer greater or equal to zero and less than the length of the string.

endIndex: the input end index. Must be a integer greater or equal to startIndex and less than the length of the string.

Return :

the new string or null if any parameter is null.

function toUpperCase(value: string): string

Returns a new string with all characters from input `value` string to upper case.

Parameters :

value: the input string

Return :

the upper cased string or null if parameter is null.

function toLowerCase(value: string): string

Returns a new string with all characters from input `value` string to lower case.

Parameters :

value: the input string

Return :

the lower cased string or null if parameter is null.

function replaceAll(value: string, pattern: string, replacement: string): string

Returns a new string with all `pattern` occurrences in input `value` string replaced with `replacement`.

Parameters :

value: the input string

pattern: the string pattern (regex) to replace

replacement: the replacement string

Return :

the new string or null if parameter is null.

function replaceFirst(value: string, pattern: string, replacement: string): string

Returns a new string with the first pattern occurrence in input `value` string replaced with replacement.

Parameters :

value: the input string

pattern: the string pattern (regex) to replace

replacement: the replacement string

Return :

the new string or null if parameter is null.

function startsWith(value: string, prefix: string): boolean

Returns `true` if a `value` string starts with a prefix.

Parameters :

value: the input string

prefix: the input prefix to look for

Return :

`true` if `value` starts with `prefix` or else `false` .

Returns `null` if any parameters is null.

function endsWith(value: string, suffix: string): boolean

Returns `true` if a `value` string ends with a suffix.

Parameters :

value: the input string

suffix: the input suffix to look for

Return :

`true` if `value` ends with `suffix` or else `false` .

Returns `null` if any parameters is null.

function contains(value: string, searchString: string): boolean

Returns `true` if a `value` string contains a `searchString` .

Parameters :

value: the input string

searchString: the input searchString to look for

Return :

`true` if `value` contains the `searchString` or else `false` .

Returns `null` if any parameters is null.

function join(separator: string ...): string

Returns a new string compose of all input substrings separated by as separator string.

Parameters :

separator: the input separator string

substring: a substring to add to the result string. All element must be string.

Return :

a new string compose of all input substrings separated by the separator input.

If any parameters is null or no substrings is provided return null.

Unit core.time

Script unit that provides time functions.

METHODS
<code>function addHours(value: time, hours: decimal): time</code> Adds hours to a time.
<code>function addMinutes(value: time, minutes: decimal): time</code> Adds minutes to a time.
<code>function addSeconds(value: time, seconds: decimal): time</code> Adds seconds to a time.
<code>function of(hours: decimal, minutes: decimal, day: decimal): time</code> Creates a time given hours, minutes and seconds.
<code>function fromDatetime(value: datetime): datetime</code> Converts a date and time to a time.
<code>function seconds(startValue: time, endValue: time): decimal</code> Returns the number of seconds between two times.

- `function addHours(value: time, hours: decimal): time`
- `function addMinutes(value: time, minutes: decimal): time`
- `function addSeconds(value: time, seconds: decimal): time`
- `function of(hours: decimal, minutes: decimal, day: decimal): time`
- `function fromDatetime(value: datetime): datetime`
- `function seconds(startValue: time, endValue: time): decimal`

function addHours(value: time, hours: decimal): time

Adds hours to a time. The calculation wraps around midnight.

Parameters :

value: the input time

hours: the number of hours to be added. Must be an integer.

Return :

the new time or null if any parameter is null.

function addMinutes(value: time, minutes: decimal): time

Adds minutes to a time. The calculation wraps around midnight.

Parameters :

value: the input time

minutes: the number of minutes to be added. Must be an integer.

Return :

the new time or null if any parameter is null.

function addSeconds(value: time, seconds: decimal): time

Adds seconds to a time. Fractions of seconds are supported with millisecond precision. The calculation wraps around midnight.

Parameters :

value: the input time

seconds: the number of seconds to be added. Fractions are rounded to three decimal digits.

Return :

the new time or null if any parameter is null.

function of(hours: decimal, minutes: decimal, day: decimal): time

Creates a time given hours, minutes and seconds. Fractions of seconds are supported with millisecond precision.

Parameters :

hour: the input hour

minute: the input minute

second: the input second. Fractions are rounded to three decimal digits.

Return :

the new time or null if any parameter is null.

function fromDatetime(value: datetime): datetime

Converts a date and time to a time. The date part is ignored.

Parameters :

value: the input date and time

Return :

the new time or null if parameter is null.

function seconds(startValue: time, endValue: time): decimal

Returns the number of seconds between two times.

Parameters :

startValue: the input start date and time inclusive

endValue: the input end date and time exclusive

Return :

the number of seconds or null if any parameter is null. Return value may have a fractional part.

Unit core.uri

Script unit that provides functions for managing Uniform Resource Identifier (URI) references.

METHODS

function of(value: string): uri Creates an URI reference from a string.

function of(value: string): uri

Creates an URI reference from a string.

Parameters :

value: the input value, for example 'https://www.tibco.com'.

Return :

the URI.

Java API Reference

- [TIBCO EBX® 6.1.0 Java API](#)

Release Notes

- [Version 6.1.0](#)
- [All release notes](#)

6.1.0

Release date: March 22, 2023.

- [Audit trail](#)
- [Tile views](#)
- [Team Up](#)
- [Hierarchies](#)
- [Record metadata](#)
- [Staging](#)
- [Workflows](#)
- [Perspectives](#)
- [Data modeling](#)
 - [Entity names](#)
 - [Search strategies](#)
- [Data validation](#)
- [Request and query optimizations](#)
- [Java API](#)
- [Data services](#)
 - [REST Toolkit APIs](#)
 - [Built-in RESTful services](#)
 - [Built-in SOAP operations](#)
- [EBX® main configuration file](#)
- [EBX® Container Edition](#)

- Limitations for this release
 - Hierarchy limitations
 - Query limitations
 - Value-labeling limitations
 - Data validation limitations
 - Linked fields limitations
- Version upgrade
 - Compilation and build environments
 - Changes on supported environments
 - Migrating an existing EBX® 5 environment
 - Upgrading from an existing EBX® 6 environment
 - Removal of dataset creation options
 - Backward compatibility
- 6.1.0 Known issues
 - User interface
- 6.1.0 Closed issues
 - Container Edition
 - Core engine
 - Data modeling
 - Data services
 - Data validation
 - Data views - hierarchies
 - Data workflows
 - Dataspaces & History
 - Documentation
 - Infrastructure
 - Queries
 - Team Up

- [User interface](#)
- [Users & Permissions](#)

Audit trail

The "Audit Trail" feature has been extended. Now it handles user access to comparison and validation services. The attribute "service-mode" is now logged in the specific case of the "User service" event. See [Audit trail](#) for more information

Tile views

You can now create a custom view that displays data using tiles. The tiles are arranged in a grid and can include the following information: title (required), subtitle, image, and description. For the title and description you can: select fields to automatically populate the values, manually specify the text that displays, or use a combination of both methods. Additionally, you can define sort and filter criteria that apply to the view. Toolbars can also be added.

Team Up

Team Up features introduce new ways to collaborate and engage with your data. Business users can leave comments on records, rate records, and assign tags to records. Access to, and visibility of Team Up features is highly customizable. See [Collaborative features \(Team Up\)](#) and [Introduction to collaborative features \(Team Up\)](#) for more information.

Hierarchies

This release contains the following new hierarchy features and updated functionality:

- In addition to Enumeration, Boolean, and Date, supported enumerable data types now also include String and Integer. See [Hierarchy limitations](#) for information on limitations for these new dependencies.
- A Group By hierarchy view can now be created based on the columns contained in a single table; the columns set the hierarchy dimensions. Fields used for this type of hierarchy must be a foreign key or an enumerable type.
- Level skipping is now available to generate views of ragged hierarchies where a member's parent skips one level. To enable ragged hierarchy display, the new property **Allow level skipping** is available under a hierarchy view configuration's **Advanced properties** group. Note

that all foreign keys must be present in the target dimension table. See [Ragged hierarchies](#) for more information.

- Fields with a Date data type can now be defined as time-based hierarchy dimensions. When reaching the last screen while creating the hierarchy view, select the date field at the top of the screen and use the **Time dimension** menu to select how you want the dimension to display. See [Time-based hierarchy dimensions](#) for more information.
- The **Remove recursive root leaves** was reactivated in this release. In a recursive case, when a search filter is applied or if **Prune hierarchy** is enabled, this option prevents root leaves from displaying.

Record metadata

Record metadata is now automatically added to a table's structure. The metadata nodes allow you to retrieve the following information about a record: creator, creation time, last modifier, time of last modification, and Unique identifier (UUID). Since this metadata is now part of the data model, you can:

- Use the quick search feature to locate records using the metadata fields as part of the search criteria.
- Execute queries with the API using XPath or SQL.
- Define table views that include the metadata fields.

See [Record metadata](#), `SchemaMetadataConstants` and `SchemaNode.getNodeChildren` for more information.

Staging

The new staging feature allows you to move elements in your EBX® repository from one environment to another. For example, you can stage configuration settings from a development environment to a testing environment. See [Staging](#) for more information.

Workflows

Several improvements were made to workflows:

Workflow Launcher	The Workflow Launcher built-in service was improved in the following ways:

	• When activating a workflow launcher you can now select and activate multiple tables at the same time. • The process to launch or view workflows that require a record selection was simplified. Now users can view and launch workflows from the toolbar that displays on rows in the table view. This prevents them from having to open a record, or use the Actions menu to locate a workflow before launching it. See Workflow launcher in data for additional information.
Workflow model editor improvements	• A pop-up is now displayed when unvalidated changes of the layout are not saved when creating a new step. • Customize the links arrangement, and more generally the steps disposition using new connectors. • A pop-up is now display when unvalidated changes of the layout are not saved. The following new automatic placement mode is available: vertical layout.
New built-in script tasks	• Update a record String value • Retrieve a record's field value (String) • Update a dataspace label • Update a workflow label

Perspectives

A new perspective mode is now available to set as the default when connecting to EBX®. Choose between the following two modes:

- **Favorite perspective:** apply the favorite perspective at connection. This is the default mode.
- **Last used perspective** (new for this release): apply the last visited perspective at connection.

See also [Perspective modes](#).

Data modeling

Data modeling includes the following new features:

- [Entity names](#)
- [Search strategies](#)

Entity names

By default, an entity name is now automatically created for tables when no name is explicitly defined in the data model. Entity names allow you to directly reference tables in SQL requests and data service operations using unique names instead of their paths. This is accomplished by defining aliases for SQL requests and suffixes for WSDL operations. This new behavior can impact existing data models.

Note that this new behavior may trigger an error on data model compilation, as indicated in the [Backward compatibility](#) section.

Search strategies

The following search strategies are now available to define on String fields: Levenshtein, Soundex, Double Metaphone, and NGram strategy. Previously, these strategies were only available when the EBX® Information Search Add-on was deployed; in case a data model was already using them, not further action is required.

See [Advanced strategies for strings](#) for more information.

Data validation

When creating a constraint on table with a record level check, you can now specify a behavior depending on the record's creation status. It is now possible to use an indicator to know if current record is being created. The method `isRecordCreation()` returns `true` if the current record is being created. See `ValueContextForValidationOnRecord.isRecordCreation` for more information.

Request and query optimizations

Several features offer better performance. In particular:

- A request cache has been introduced that caches the transformation from requests to queries, when executing a `Request`. See [Request cache activation](#) for more information.
- `RequestResult` size related methods have a better performance if executed multiple times.
- Requests that are based on primary key values will be executed using a more optimized implementation of `RequestResult`.

Java API

The Java API was improved on the following topics:

Iterable <code>RequestResult</code>	<code>RequestResult</code> now extends <code>Iterable<Adaptation></code> so it can be traversed using a simplified syntax. See <code>RequestResult</code> for more information.
<code>Repository</code>	The new method <code>Repository.getDefaultWhenFullyInitialized()</code> blocks the calling thread until the repository is fully initialized and become able to handle: SOAP, REST Built-in, REST Toolkit requests and UI. See <code>Repository.getDefaultWhenFullyInitialized</code> for more information.

<code>SchemaLocation</code>	Some methods have been added to <code>SchemaLocation</code> to retrieve the properties of packaged and embedded data models. See <code>SchemaLocation.isPublication</code>, <code>SchemaLocation.getPublicationName</code>, and <code>SchemaLocation.getPathInModule</code> for more information.
<code>UICustomViewContext</code>	The new Java interface <code>UICustomViewContext</code> allows to get: • the recommended view for a given user; • the favorite view for a given user; • the currently applied view. See <code>UICustomViewContext</code> for more information.
<code>ValidationReportItemSubject</code>	Two new methods have been added to <code>ValidationReportItemSubject</code> to check the subject type of a validation item. • See <code>ValidationReportItemSubject.isDatasetOrRecord</code> for more information. • See <code>ValidationReportItemSubject.isTableSubject</code> for more information.

Data services

The following sections describe new features and improvements to data services:

- [REST Toolkit APIs](#)
- [Built-in RESTful services](#)
- [Built-in SOAP operations](#)

REST Toolkit APIs

The OpenAPI operations now extend the 'api' category to generation documents for REST Toolkit resources. See [OpenAPI](#) for more information.

Built-in RESTful services

The built-in RESTful services were improved in the following areas:

URL encoding	URL encoding has been improved for REST services. It now complies with the following URI specifications: RFC-3986 with percent encoding (section 2.1) . See backward compatibility for more information.
Record metadata	The select operation has been improved to return record metadata using the new <code>includeMetadata</code> parameter. The system metadata can now be used by the <code>filter</code> and <code>sort</code> parameters with the following metadata fields: creator, creation time, updater and update time. As a consequence the: • <code>includeMeta</code> is deprecated and replaced by the <code>includeMetamodel</code> parameter. • <code>includeTechnicals</code> parameter is deprecated and replaced by the <code>includeMetadata</code> parameter with the 'system' value. • <code>checkNotChangedSinceLastUpdateDate</code> parameter is deprecated and replaced by the

	<code>checkNotChangedSinceLastUpdateTime</code> parameter for the delete operation. Update and delete operations now support the Optimistic locking in JSON compact format with the same mechanism as the extended format. See data modeling Record metadata, <code>includeMetadata</code> parameter, <code>includeMetamodel</code> parameter and <code>checkNotChangedSinceLastUpdateTime</code> parameter for more information.
New mass update and delete operations	These operations can combine a XPath predicate expression and/or a published table view that can also contain an XPath expression. The update operation applies one or more changes to each record that matches. See Insert operation and Delete operation for more information.
Health REST built-in	New health category operation helps monitor the EBX® server. See Health operations for more information.
Lightweight JSON extended format	The selector URL of the multi-occurs enumeration field is moved under the field instead of each occurrence. This change lightens the JSON content for select operations.

Built-in SOAP operations

The built-in SOAP operations were improved in the following areas:

Record metadata

The select operation was improved to return record metadata using the new `includesMetadata` parameter. The system metadata can now be used by the predicate parameter with those fields: creator, creation time, updater and update time. As a consequence, the `includesTechnicalData` is deprecated and replaced by the `includesMetadata` parameter with the 'system' value.

The update operation was improved to take all data content returned by the select operation even if fields are ignored. For example, it can take association values and record metadata, except `ebx-metadata/system/update_time` to ensure success of the [optimistic locking](#) feature.

See `includesMetadata` [parameter](#), [Update operation](#) and [Record metadata](#) for more information.

EBX® main configuration file

The main configuration file has been improved on the following topics:

Configuring SMTP without JNDI

Timeouts for sending an email are now configurable from main configuration file, when ignoring the Java Naming and Directory Interface (JNDI).

See [Activating and configuring SMTP and emails](#) for more information.

EBX® Container Edition

The Container Edition was improved on the following topics:

Java SE Development Kit	The JDK has been upgraded to LTS version 17. See Overview for more information.
Configuring SMTP timeouts	Timeouts for sending an email are now configurable. See Email connectivity for more information.

Limitations for this release

Hierarchy limitations

On allowed dimension for new dependencies:

- List field is not allowed inside a Nested Column Dependency.
- Date List field is not allowed as a Time Dimension field.

Query limitations

The performance of these features is not yet optimized:

- Navigation on paginated views (next/previous page) with sort on foreign key fields having a custom display pattern, if this pattern contains fields outside the primary key of the referenced table.
- Advanced search filtering on foreign key fields using string-based operations (e.g. starts-with, contains, matches, etc.).
- Navigation on paginated views (next/previous page) with sort on linked fields.

The following cases do not rely on Lucene indices. This has two consequences: first, performance is not optimized; second, Quick search behaves like a 'contains', rather than a full-text search:

- computed fields with non-local dependency;
- inherited fields;
- all fields in an inherited dataset;
- all fields in table history views.

More specific cases not benefiting from index optimization are described in the [following section](#).

Value-labeling limitations

Search and sort using the value-labeling feature have a few limitations described in section [Value-labeling](#).

Data validation limitations

- It is not possible to validate a dataset or a table within a Procedure.

Linked fields limitations

- Targeting a node which is an enumeration is not fully supported yet. Currently, only the raw value will be available for display and search, instead of the enumeration label. Until this limitation is removed, the workaround is to declare the same enumeration in the target node and in the linked field itself.

Version upgrade

In addition to the information provided in this release note, the chapter [6.1 Java API changes](#) provides an exhaustive list of the new Java APIs published since the release of the 6.0 version.

Compilation and build environments

Your compilation and build environments must now use the following new development JAR libraries, instead of `ebx.jar`:

- `ebx-api.jar`, for the Java API of EBX®;
- `ebx-resources-api.jar`, for the Java API of third parties.

Note:

Java types in `ebx.jar` are no longer obfuscated.

Note:

At runtime, development JAR libraries should not be deployed, continue to only use the `ebx.jar` library.

Changes on supported environments

The chapter [Supported environments](#) has been updated to reflect the necessary upgrades detailed below.

Browser support	Please note that Microsoft Internet Explorer and the non-Chromium version of Microsoft Edge are not supported anymore.
JVM support	Required Java Runtime Environment: either JRE 11 or 17. Java 8 is no longer supported.
IBM support	WebSphere Application Server 9 is no longer supported. WebSphere Application Server Liberty, for Java EE 8 only, is supported.
Oracle support	Oracle 12c is no longer supported. Minimum supported version of Oracle is now 19c.
PostgreSQL support	PostgreSQL 10 is no longer supported. Minimum supported version of PostgreSQL is now 11.
Microsoft support	SQL Server 2012 SP4 is no longer supported. Minimum supported version of SQL Server is now 2014.

Migrating an existing EBX® 5 environment

To get a running EBX® 6 from an existing EBX® 5 environment, follow the [dedicated migration guide](#).

Upgrading from an existing EBX® 6 environment

The Lucene library version has been bumped from 8.11.0 to 9.4.0, bringing modifications in the index format. As a consequence, the directory holding the indices is going to be automatically deleted during the product upgrade. Thus, a delay at the first access to each table can be expected for index rebuild. Some minor changes in search behavior can also be observed.

Removal of dataset creation options

In the **Create a dataset** service, the option to create a dataset from an *external data model* was removed. It is still possible to import a file in the data model assistant.

Backward compatibility

URL encoding	URL encoding has been improved for REST services and GUI. It now complies with the following URI specifications: RFC-3986 with percent encoding (section 2.1) . The percent encoding is used to encode path and query parameters values and keys of URL when they contain dataspace keys or primary keys which may have characters that would require encoding. For example, '&' in a query parameter value of a URL (like in <code>A&B</code>) will be encoded as <code>A%26B</code>) Server configuration should be checked to support slash and backslash characters. See Installation notes for more information.
Entity names	In this release EBX® now automatically creates an entity name for the tables defined in a data model. The last step of the table's path is used as the default entity name. Since entity names must be unique in a data model, an error is now raised during compilation if the model contains tables with the same name. To prevent an error in this

	case, you must explicitly set the entity names. See Entity mappings and <code>SchemaNode.getTableEntityName</code> for more information.
JSON extended format change	Selector URL has been moved for multi occurrence enumeration fields using the select operation. See Lightweight JSON extended format for more information.
Maximum number of clauses in a query	Due to the upgrade of the Lucene library, the method to compute the maximum number of clauses authorized in a query has changed. Previously, it was possible to have a very high number of clauses. Those queries now throw an exception. The condition is not easily met. If it happens, this is probably due the Levenshtein search strategy. This can be avoided by customizing the strategy to decrease its fuzziness. For instance, by adding an invariant prefix, or reduce the max edit distance.
Rolling log files	From 6.0.6 to 6.0.10 (included), the rolling files names have changed from <code>xxxx.log.yyyy-MM-dd</code> (or <code>xxxx.log.yyyy-MM-dd-i</code>) to <code>xxxx.MM-dd-yyyy.log</code> (or <code>xxxx.MM-dd-yyyy-i.log</code>). By default, the former format (<code>xxxx.log.yyyy-MM-dd</code> or <code>xxxx.log.yyyy-MM-dd-i</code>) is used. The rolling format can be specified in <code>ebx.properties</code> file by setting the property <code>ebx.log4j.appender.ebxFile.rollingFormat</code>. If this property is specified, the filenames will have the following format:

`xxxxx.rollingFormat.log` and will not be append to the end of the filename with the extension. To configure same filename format than since 6.0.6, the value must be `%d{MM-dd-yyyy}-%i`. The property will accept both a date/time pattern compatible with [Java SimpleDateFormat](#) and/or a `%i` which represents an integer counter. The integer counter is used in conjunction of the backup threshold, and will increment to create new files when the current file size will reach the threshold.

Since 6.0.6, the

`ebx.log4j.appender.ebxFile.backup.Threshold` property defines the maximum size (in megabytes) of backup files for daily rollover and *not* the maximum number of backup files.

6.1.0 Known issues

User interface

- **CP-24388** Tiles views cannot be used as advanced selector of a combo box.

If a tile view is specified as the advanced selector of a combo box, the "select" button is not displayed in the cards of the tile view.

6.1.0 Closed issues

Container Edition

- **CP-24408** A stopped EBX Container Edition container may fail when started again.

Core engine

- **CP-18256** Sorting a table with large volume on a foreign key is very long.

Sorting a table with a large volume of data takes longer than expected.

- **CP-18919** EBX logs are not deleted when the file size is zero.
- **CP-20968** The loading screen displays for a long time when applying the `<emphasis role="bold">Apply last modifications</emphasis>` service.

This occurs when the service is run on a child dataset that contains 100,000 records.

- **CP-21133** The Jaro-Winkler strategy should not accept parameters.
- **CP-23002** A `NullPointerException` occurs when applying a Jaro-Winkler search for Assigned search strategies.

The exception happens when executing a global search and the Assigned search strategies for a field is set to `Default search template` and uses the Jaro-Winkler strategy.

- **CP-23572** The "present in" operator is not available on an enumeration field of the HTML data type.
- **CP-24178** `AdaptationHome.getTerminationDate()` does not return null on open dataspace.
- **CP-24189** Clustering throws errors when relation matching and cluster field in complex data type.

Clustering throws errors when using relationships in matching and the cluster field is in a complex data type.

- **CP-24199** Datasets cannot be accessed during data model compilation.
- **CP-24257** A `NullPointerException` occurs when levels for logging categories are defined by a variable.
- **CP-24290** Error importing a perspective archive when a perspective has been duplicated.

An error occurs when importing an archive that contains a perspective that was duplicated.

- **CP-24333** The add-on license threads can block the server shutdown.
- **CP-24345** `NullPointerException` when exporting a record, table or dataset created in severe error.

A `NullPointerException` occurs when exporting a record, table, or dataset that was created in a state of severe error.

- **CP-24346** `ClassCastException` when calling `SchemaFacetTableRef.getLinkedRecords()` in a Widget

A `ClassCastException` occurs when calling `SchemaFacetTableRef.getLinkedRecords()` in a widget.

- **CP-24379** Repository purge operation fails at validation reports cleaning step.

A repository purge operation fails at the validation report cleaning step.

- **CP-24405** `UIDDropDownList` on a foreign key does not display elements.
- **CP-24449** Migration to EBX 6 fails if the metadata references tables that don't actually exist.

Migration to EBX 6 fails if the metadata references tables that don't actually exist.

- **CP-24579** An error occurs when rebuilding indexes in a child dataspace.
- **CP-24627** Default value is set on existing records in history tables when adding a new column

When adding a new column in a historized table its default value is added to the existing records in the history.

- **CP-24634** Errors prevent access of table contents.

This includes errors with messages like the following: `java.lang.IllegalStateException: Cf0ccurrence with ref (...,...) could not be fetched from table zza7[<tableName> of content type=VALUES in dataset <datasetName> in dataspace DataspaceName]: record have been deleted a transaction`.

- **CP-24657** "java.lang.IllegalStateException: This index is read-only" error while updating a record.

An `IllegalStateException` error can occur when updating a record.

- **CP-24706** Editing Greatest Value set field to a greater value doesn't affect Auto-increment when creating a new record.
- **CP-24733** Sub queries are not correctly identified when preparing query parameters.
- **CP-24763** Missing `com.orchestranetworks.service.directory.EncryptPassword` class.

The `com.orchestranetworks.service.directory.EncryptPassword` class mentioned in the main configuration file documentation is missing in the product package.

- **CP-24895** Values from uncommitted transactions can corrupt table size estimation.
- **CP-24917** An error occurs when applying an advanced search on the following table: **Administration > Database mapping dataset > Tables**.
- **CP-24944** When extending `SchemaDocumentation`, displayed links are not clickable.
- **CP-24997** There is a missing search strategy for the default template.

An exception can be thrown when compiling a data model that contains a foreign key node with multiple search strategies defined.

- **CP-24999** A `LockObtainFailedException: Lock held by this virtual machine` error occurs.

If something goes wrong while rebuilding an index, the `IndexWriter` used to perform the rebuild is not properly closed and the slot remains locked preventing it from being reused afterwards.

- **CP-25032** Log failure can occur during index refresh at commit.

This can disrupt synchronization between the cached index and the table content after commit.

- **CP-25233** A primary key violation occurs when replicating records with lists in a mass update.
- **CP-25270** A `LxiDuplicateKeyException` occurs when creating a record.
- **CP-25405** When the `XPxParametrizedLiteral` value is in an internal format, it can cause a `ClassCastException` to occur inside of `toXPathExpression`.

Data modeling

- **CP-21255** SqlChecker error with data models if they have circular dependencies, including one data model in error
- **CP-22225** Deprecated properties of table reference cannot be used

The data model does not compile, when declaring a table reference using the deprecated properties: `displayKey` or `labelPaths`.

Data services

- **CP-2130** REST built-in select does not return data from table that has multiple-occurrences group.

A REST request does not return data from the inner group when a multiple-occurrence group is nested within a multiple-occurrence group.

- **CP-17619** Optional parameters in SOAP header do not work using default namespace without prefix.

When calling a SOAP operation, adding additional parameters in the header does not work when using the default namespace without prefix.

- **CP-20709** `ContentHolderForInput` automatically removes object attributes named "content".

When a `ContentHolder` is used for data input, it automatically removes object attributes named "content" even if these attributes are parts of the data model.

- **CP-21814** REST count operation on association field with `:count` syntax returns 404.
- **CP-21830** The documentation does not mention the limitation for procedure context with the REST built-in JSON compact format
- **CP-24606** Lineage WSDL configuration does not considered UI selected context.

The lineage WSDL configuration does take the selected UI context into account.

- **CP-24653** The REST built-in transient operation should ignore association fields and value function nodes.
- **CP-24664** SOAP body validation error for select operation with some activated parameters.

When a select operation enables `exportCredentials` with a value of `true`, the XML body response is invalid.

- **CP-24673** NPE when branch does not exists using SOAP operation.

When using a SOAP operation, a `NullPointerException` occurs if the defined branch does not exist.

- **CP-24722** Association field must be ignored using REST builtin import/update operations.

When using the REST built-in import/update operations, ignore association fields.

- **CP-24962** A `PathAccessException` occurs when importing XML.

The exception causes the import to fail when a record contains an empty group that comes directly after a list of groups.

- **CP-25013** A REST built-in selector field must not be present for a read-only field.

Data validation

- **CP-7461** Cannot enter a date with format like "dd/MMM/yyyy"

It is not possible to enter a date if the formatting policy specifies a 3-character format for the month.

- **CP-18807** In a user service, an exception may be thrown on page redisplay.

The User Service API can be used to customize the form for creating a record. If a record contains a foreign key, an unexpected exception may be thrown, when the user service is redisplayed.

- **CP-21584** In a `UserService`, the record form submission is not blocked when an error message is added in `ConstraintOnTableWithRecordLevelCheck#checkRecord`.
- **CP-23502** The record validation report is computed for each tab to be displayed in a record form

The header of a record's form displays the up-to-date validation report of the related record. Each tab of the form has its own header. As a consequence the validation report is computed as many times as the number of tabs to be displayed in the form. This can lead to performance issues if the table contains many constraints.

- **CP-24171** `CalciteException` Table not found when validating a table with a table reference filter using a table reference node

An exception is thrown when a table reference filter uses a node that is a table reference to a table located in another dataset.

- **CP-24192** `NullPointerException` in code generated by calcite for a table reference filter validation

A `NullPointerException` occurs in code generated by Calcite for a table reference filter validation.

- **CP-24336** `xs:pattern` facet is not applied on the XML string representation of dates.

The `xs:pattern` facet checks the string representation of the Java Date object instead of the XML string representation.

- **CP-24483** `NullPointerException` while updating the persisted validation report

A `NullPointerException` occurs while updating the persisted validation report.

- **CP-24526** Slow performance when accessing the persisted validation report.

Performance is impacted when accessing the persisted validation report.

- **CP-24777** An `IllegalStateException` occurs when accessing a validation dataspace after an aborted deletion of a functional dataset.

- **CP-24904** Validation of a table reference fails with a `ClassCastException`.

This is caused when `PvrActionForTableRef` generates a query plan with the wrong `FK_AS_STRING` parameters

- **CP-24923** A `ClassCastException` occurs during the validation of a contextual table reference filter.

- **CP-24933** A contextual table reference filter with multi-valued fields returns incorrect results.
- **CP-25019** The result of `ValidationReportItemSubjectForTable.getRecords()` contains `null` values.

Data views - hierarchies

- **CP-16830** Hierarchy cycle detection does not work on a recursive ternary dependency.
- **CP-17169** Cannot execute the specific property "prune recursive orphan leaves".

The option to prune orphan leaves cannot be executed if there are recursive members without a parent.

- **CP-17204** The node level icons and their text are not aligned in the hierarchy view creation screen.

The node level icons and text are not aligned in the hierarchy view creation screen.

- **CP-18439** Search not found on a multi links hierarchy.

Search results are not found on a multi-link hierarchy.

- **CP-20401** Some child nodes are duplicated in a hierarchy view.

When a hierarchy view is applied, some child node display twice.

- **CP-20448** Level filter on enumerable does not work for hierarchy.

An incorrect result is shown when users create a hierarchy view and set a filter.

- **CP-21838** The hierarchy 'expand all' is not anymore limited to 2000 displayed members.

Hierarchy expansion is no longer limited to displaying 2000 members.

- **CP-21840** Find on a ternary recursive hierarchy fails.

The Find tool on a ternary recursive dimension hierarchy does not work properly.

- **CP-21882** The 'Last' button is activated on the last page of a hierarchy view based on the Date/ enumeration field.

After navigating to the last page in a hierarchy that is based on a field with a data type of date or enumeration, the "move to last" navigation button is still active.

- **CP-21892** `member.isLooping()` may lead to `stackOverflow` mostly in workflow.

Use of `member.isLooping()` might lead to a `StackOverflowError`, mostly in workflows.

- **CP-21900** The child nodes of levels are not displayed fully when creating a hierarchy view.

When users create and apply a hierarchy view that has several levels, its child nodes are not displayed fully.

- **CP-22403** Number of selected records is displayed incorrectly on toolbar when users select/unselect recursive record on hierarchy view multiple times.

The number of selected records does not display correctly in a hierarch view toolbar when a recursive record is selected and deselected multiple times.

- **CP-22429** Nodes of a hierarchy view on a ternary with an ordering field model are switched after enabling pruned mode.

The order of nodes is switched after enabling pruned mode on a ternary table hierarchy view.

- **CP-22483** In a hierarchical view, some records are displayed although they should be filtered out.

Some records display in a hierarchical view when they should be filtered out.

- **CP-22881** Hierarchy branch selection does not work on ragged hierarchy.

Hierarchy branch selection does not work on a ragged hierarchy.

- **CP-23010** An unexpected node is displayed on a hierarchy view.

An unexpected node displays in hierarchy views and an error occurs if it is selected.

- **CP-24281** Reimplement the 'prune recursive orphan leaves' option.

The **Remove recursive root leaves** option must be reimplemented.

- **CP-24443** 'Enable/Disable pruned mode' is not properly set.

The pruned mode options are not set properly.

- **CP-24558** Pagination does not work on an enumerable field.
- **CP-24646** Hidden records must not be used for hierarchy display.
- **CP-24675** The cycle member remains, in a ternary dependency hierarchy.

In a ternary dependency hierarchy containing a cycle member, the cycle member remains alone after a search; it is not in a valid branch.

- **CP-24746** Hierarchy REST API orphan members are not properly managed.

In the REST API, a hierarchy's orphan members are not managed properly.

Data workflows

- **CP-19283** Documentation in activation record is not updated when updating the label of a Launcher.

A launcher activation label can be inherited from the label of its launcher or customized. If this label is inherited, it is not updated when the label of the launcher is modified and users might not realize this. No information or hint displays to inform users that they must save again to apply new values.

- **CP-22631** Fix "remove" icon in workflow diagram editor.

The **Remove** icon in the workflow diagram editor does not display correctly.

- **CP-24148** No way to clear a profile field in a workflow step

There is no way to clear a profile field in a workflow step.

- **CP-24578** Workflow complex objects should be expanded by default.
- **CP-24724** When dragging a link to another step and cancelling the confirmation, the link end spot changes.
- **CP-24769** The documentation for the Xpath parameter of the update a record field value script task is not correct.
- **CP-25239** Data context mapping in the model configuration should be open by default.

Dataspaces & History

- **CP-24934** A `NullPointerException` occurs during the dataspace compare action.

Documentation

- **CP-18443** Fix java sample in security documentation.

The Java sample in the documentation on security best practices needs to be updated.

- **CP-25026** Incorrect Java EE platform for WebSphere AS 8.5 and 9.0 support

The Java EE platform for WebSphere Application Server 8.5 and 9.0 support is incorrect.

- **CP-25092** The double quote usage for the Windows path in the Tomcat installation note is incorrect.

Infrastructure

- **CP-24586** All `handleRepositoryStartup` method should complete before opening access to EBX®.

Queries

- **CP-24291** Field trimmer removes all fields from the input to a LogicalAggregate with COUNT().

A `ClassCastException` occurs during the field trimming step.

- **CP-24319** Result of `RequestResult.getSize()` might be incorrect and vary from call to call.

The `RequestResult.getSize()` method might return an incorrect result and vary from call to call.

- **CP-24551** Performance problem when setting request log level to DEBUG because of a slow `RelNode` to SQL translation.

A performance issue occurs when setting the request log level to DEBUG.

- **CP-24608** Although set as non confidential, hidden nodes cannot be used in a foreign key filter when under a complex node.
- **CP-24631** Sort by resource field (and pagination) when the Lucene index is not available throws an exception.
- **CP-24677** An SQL query throws a `CannotPlanException` due to `HepRelVertex` in the plan after the Hep phase.

Team Up

- **CP-24911** An error occurs when applying search with Team up fields on hierarchy view.

User interface

- **CP-2041** Fields are unexpectedly filled automatically while creating a record through an association.

When creating a record through an association, fields in the creation form should be pre-filled, but not read-only.

- **CP-17509** An error occurs when using a web component URL that selects a perspective and a branch.

- **CP-19856** The page for selecting backed-up user does not display in some cases

A user selection page backup does not display when defining a custom login.

- **CP-20087** An alignment problem occurs between the text of button and its arrow.

On the **Import** screen and **Task Creation** modal, the button text and arrow do not align horizontally.

- **CP-20157** The toolbar overlaps the number of filters applied when searching on a table under data workflow.

When users take a task and apply an advanced search, the toolbar overlaps the number of filters applied.

- **CP-21620** [WCAG 2.0 - 4.1.2](#). The 'Remove' and 'Checkbox' buttons should provide additional information when using the accessibility checker tools (NVDA).

Additional information should include its role, state, and value related information such as the element removed and checked, and it's corresponding value.

- **CP-21684** The tables under a group still display when the group is set to hidden.
- **CP-21753** An alignment problem occurs in the Wait task creation screen of Workflow models.

In the workflow model's **Wait task** creation screen, the "x" icon for the **Hidden in progress view** fields are not vertically aligned with the radio buttons **Yes** and **No**.

- **CP-22032** A corresponding dataset is not selected when selecting the 'Data' column on the 'Export' dataspace screen.

The dataset selection behavior does not work correctly on the **Export** dataspace screen.

- **CP-22041** The URL computing documentation must mention the data lineage configuration when using a reverse proxy.

When using a reverse proxy (particularly in https), it is necessary to configure the data lineage property. This should be mentioned in the URL computing configuration.

- **CP-22275** Default values are lost if a custom widget is defined.

In certain circumstances default values are lost when a custom widget is defined.

- **CP-22295** The scrollbar is not displayed on 'Refresh all data models'.

When invoking 'Refresh all data models', the scrollbar is not displayed in the pop-up.

- **CP-22703** Increase advanced search pane height produces a bad display.

Display issues occur in the hierarchical and tile views when re-sizing the advanced search pane.

- **CP-22970** The default label is missing in the Workflow launcher records.
- **CP-23788** When applying an advanced search with the "present in" operator, the value of the drop-down list field is duplicated sometimes.

The value of the drop-down list field is duplicated when applying an advanced search with the "present in" operator.

- **CP-24131** The highlighted bar is hidden in Data Model and Workflow sections.
- **CP-24668** A confirmation message is not displayed for a user service declared on a dataset in a perspective.
- **CP-24771** The bottom horizontal scroll bar doesn't display when loading tables in the UI.
- **CP-25000** Errors are not clickable.
- **CP-25166** When the label of a record is empty, the breadcrumb and its documentation pane should display a placeholder.

Users & Permissions

- **CP-21226** Exception when filtering on linked field if the FK is hidden.

An exception is thrown when filtering on a linked field if the foreign key is hidden.

- **CP-24489** The dropdown to set the owner of perspectives is only showing profiles with the built-in administrator role.