



TIBCO EBX®

Release Notes

Version 6.2.0
June 2024

Table of contents

| | |
|----------------------------------------------------|----|
| Workflows..... | 5 |
| Performance..... | 5 |
| Data validation..... | 5 |
| Application servers..... | 6 |
| Data services..... | 6 |
| Built-in RESTful services..... | 7 |
| Built-in LDAP directory..... | 7 |
| Improving the performance with prefilters..... | 7 |
| EBX® Container Edition..... | 7 |
| EBX® Script IDE..... | 8 |
| Home page..... | 8 |
| Metadata management..... | 8 |
| History view of a record on the sidebar..... | 9 |
| Product migration..... | 9 |
| Limitations for this release..... | 9 |
| Hierarchy limitations..... | 9 |
| Query limitations..... | 9 |
| Value-labeling limitations..... | 10 |
| Data validation limitations..... | 10 |
| Version upgrade..... | 10 |
| New Web Application Resource (WAR) files..... | 10 |
| Changes on supported environments..... | 10 |
| Upgrading from an existing EBX® 6 environment..... | 11 |
| Backward compatibility..... | 12 |
| 6.2.0 Known issues..... | 13 |
| Data Services..... | 13 |
| User Interface..... | 13 |
| 6.2.0 Closed issues..... | 14 |
| Container Edition..... | 14 |
| Core Engine..... | 14 |
| Data Modeling..... | 15 |
| Data Services..... | 15 |
| Data Validation..... | 15 |
| Data Workflows..... | 15 |
| Hierarchies..... | 15 |
| Infrastructure..... | 15 |
| Scripting..... | 15 |
| Team Up..... | 15 |
| User Interface..... | 15 |
| Users & Permissions..... | 16 |

Release date: June 21, 2024.

Workflows

Notable points done for workflows :

Workflow history

Fixed the service "clear from workflows selected" and "Clean from a date" available in Workflow history administration. Only records in workflow table were cleared since 6.0, some related data to it were ignored causing the history to grow bigger without the knowledge of the users. We still recommend to regularly launch a standard purge task as mentioned in workflow administration documentation.

See also *Workflow Administration*

Performance

A prefilter can now be defined for an `AdaptationFilterAPI` or `TableRefFilterAPI` to improve query performance or table reference validation.

See their Java documentation or the [Improving the performance with prefilters](#) section for more information.

Data validation

This service pack improves the performance of creating a child dataspace from a dataspace that contains many datasets with validation reports. As part of this improvement, the validation report metadata schema was changed. As a result, all validation reports will be reset when upgrading from EBX® 6.0, or 6.1 to 6.2 or later. The stale data in the reset validation reports will be removed during the next purge.

Application servers

This section describes application servers changes for architecture modernization.

Servlet and JSP container

EBX® now supports Jakarta EE 9 or above. This change will require an upgrade to several third party libraries, such as Servlet 5.0 or greater, JSP 3.0 or greater, JMS 3.0 or greater, and Mail 2.0 or greater. See the links below for more information on requirements for specific components.

Developers of EBX® 6.2.x onwards should be aware that, as a result of the move from Java EE to Jakarta EE, the primary package name for all implemented APIs has changed from `javax.*` to `jakarta.*`. This will almost certainly require code changes to enable modules or other libraries to migrate from EBX® 6.1 and earlier to EBX® 6.2 and later.

See *Supported application servers*, *Java Message Service (JMS)*, *JMS for distributed data delivery (D3)* and *Installation notes* for more information.

Java Compile time Environment

EBX® sources are now compiled with JDK 17. JRE 11 is no longer compatible. We recommend upgrading your JRE to a recent LTS version to take advantage of its performance improvements.

See *Supported application servers* for more information.

Data services

The following sections describe new feature and improvement to data services.

Built-in RESTful services

The built-in RESTful services were improved in the following area:

Business Object

Improves select business object operation using data-bo and data-compact-bo and now returns content of one to many relationship of business objects. As the response may be large, the maximum size can be overloaded on the EBX® server using the `ebx.dataservices.rest.bo.maxResponseSizeInKB` property. In fact, if the response approaches or exceeds the max size value a HTTP 206 Partial Content response is returned with a number of records lower than requested, including a consistent paging context, all are in a well-formed JSON response.

See [ebx.dataservices.rest.bo.maxResponseSizeInKB, backward compatibility for built-in RESTful services, Performance and Select operation](#) for more information.

Built-in LDAP directory

A built-in support for LDAP (Lightweight Directory Access Protocol) has been added and designed as a custom directory implementation. This feature enables integration of an existing LDAP directory with EBX®.

See *Built-in LDAP directory* and *Built-in LDAP directory* properties, for more information.

Improving the performance with prefilters

Programmatic filters can include a combination of simple checks and more complex Java code. Under certain conditions, the simple checks can be converted to a prefilter in the form of an XPath predicate. The prefilters can then be used by EBX® to optimize queries. The prefilter must satisfy a consistency condition and it should satisfy certain other conditions for performance reasons. See the Java documentation of the relevant functions.

The interface `PrefilteredAdaptationFilterAPI` has been added for adding a prefilter to `anAdaptationFilterAPI`.

A `TableRefFilterAPI` can define an XPath prefilter during the setup of the table reference filter. See `TableRefFilterContext.setPrefilterAPI`.

EBX® Container Edition

The Container Edition was improved on the following topic:

Java SE Development Kit

The JDK has been upgraded to LTS version 21.

See *Overview* for more information.

EBX® Script IDE

This release introduces the Script IDE, a web-based development environment that allows you to create and manage scripts written in the EBX® Script Language (ESL). The web-based IDE simplifies scripting tasks such as writing functions and procedures applied to basic data manipulation, automated triggers, and even complex workflow tasks.

The Script IDE helps avoid the overhead associated with writing, compiling, and deploying custom modules. The combination of ESL and the new Script IDE make programmatic data management tasks more accessible to those without a Java programming background.

The Script IDE provides many of the features you would expect to find in an IDE, such as:

- Syntax Highlighting
- Contextual Code Completion
- Error Detection
- Version History and Rollback

See also

Overview of the Script IDE

Script reference documentation

Home page

This release introduces the home page, a new customizable entry point to EBX®.

See also *Home page*

Metadata management

This release introduces support for metadata management, offering a suite of new features:

- Data Catalog: Organize and improve findability of data across your organization.
- Business Glossary: Define business terms and link them directly to your data assets.
- Metadata-Specific Views: Gain specialized perspectives on your metadata for better analysis and reporting.
- Governance Functionality: Implement data governance policies to maintain data integrity and compliance.

Additionally, you can connect to and integrate with a variety of external data sources to extract, discover, and manage your metadata. These enhancements aim to improve your ability to manage data efficiently and make informed decisions based on reliable, well-governed data.

Business user and administrative features are accessible from two pre-defined perspectives: **Metadata** and **Metadata administration**. These built-in perspectives are available *only* if the metadata module is deployed.

Attention

As of this release, metadata management features are only available for the TIBCO EBX® Container Edition.

See also

Metadata introduction

Administration perspective

History view of a record on the sidebar

Users can now view record transaction history on the sidebar. This view shows a list of transactions, modified field values, the record modifier, and timestamp. Transactions are sortable by date and time.

Product migration

This release contains an updated *Migration Guide* that provides instructions for migrating from EBX® version 5.9.x to 6.x.

Limitations for this release

Hierarchy limitations

On allowed dimension for new dependencies:

- Recursive list field is not allowed.
- List field is not allowed inside a Group List.
- List field is not allowed inside a Nested Column Dependency.
- Recursive behavior 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.

Version upgrade

In addition to the information provided in this release note, the chapter 6.2 Java API changes provides an exhaustive list of the new Java APIs published since the release of the 6.1 version.

New Web Application Resource (WAR) files

| | |
|--------------------|-------------------------------------------------------------------------------------------------------------------|
| ebx-ide.war | This is the WAR for the new Integrated Development Environment. <i>See also Overview of the Script IDE</i> |
| <hr/> | |
| ebx-hub.war | This is the WAR for the new Hub web application. <i>See also Home page</i> |

Changes on supported environments

The chapter *Supported environments* has been updated to reflect the necessary upgrades detailed below.

| | |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hana DB support | Support for SAP HANA as a back-end database engine for EBX® is discontinued as of this release. |
| <hr/> | |
| WebLogic AS support | Currently, no version of WebLogic AS supports JDK 17, or 21 LTS and Jakarta EE 9. Please be aware that, according to Oracle's Statement of Direction, future versions of supported software might not be compatible with EBX®. |

Upgrading from an existing EBX® 6 environment

The internal versioning mechanism of indices has been improved and the Lucene library version has been bumped to 9.9.1, causing modifications to 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 to rebuild the indices. Some minor changes in search behavior can also be observed.

The upgrade process crawls all the records in the G_BLK database table to ensure their size is optimal by removing potential wasted space and splitting entries that became bigger than necessary. This one-shot process may significantly delay the first startup on large repositories.

Backward compatibility

| | |
|--------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nested procedure use of its parent's ProcedureContext | <p>A nested procedure now cannot use the <code>ProcedureContext</code> of its parent procedure to alter the parent's dataspace. Modifications made by the nested procedure are confined to its own dataspace. This restriction is in place due to concerns about data consistency, as changes made in the nested procedure might rely on unfinalized or cancelled changes in the parent procedure.</p> <p>Attempts to operate on the <code>ProcedureContext</code> of the parent procedure will now result in an <code>IllegalStateException</code>.</p> |
| Data validation | <p>The validation reports created by EBX® versions prior to 6.2 are not compatible and will be reset. The stale data for the reset validation reports will be removed during the next purge.</p> |
| Optimization of matches predicate | <p>As a consequence of the bug fix CP-26571 related to the processing of matches and matches-case-insensitive predicates, certain operations with these predicates are no longer optimized using Lucene. The reason for this is that Lucene's regular expression engine does not support certain Java regex constructs, such as <code>^</code> (beginning of line) or <code>\$</code> (end of line). Predicates using these operators were incorrectly optimized and returned incorrect results in previous versions. Therefore if the matches pattern contains such unsupported features, or if it cannot be determined (e.g. the pattern is a dynamic parameter whose actual value is unknown at the optimization phase), then the operation cannot be optimized using Lucene and the performance can be impacted.</p> |
| Association with inconsistent definition | <p>If an association is not consistent in the context of a specific dataset then it is now considered as read-only from a permissions standpoint, instead of being hidden as in previous versions. The association is read-only in the user interface and it is not possible to create, delete, attach, detach or move an associated record. It is possible to modify an existing associated record.</p> <p>As a consequence of this change an exception is no longer raised in the existing API that compute the result of the association.</p> <p>See <code>AssociationLink.getAssociationResult^{API}</code> and <code>AssociationLink.getTargetTable^{API}</code></p> <p>For more information, see <i>Association</i>.</p> |

| | |
|------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Built-in RESTful services | In business objects selection, the one to many relationship is no longer navigable through HATEOAS links, but is directly included in the JSON response. |
| Redirection to HTTPS is based on a new option | <p>Since EBX® 6.1.3 HF-1, when users accessed the login screen using HTTP and the property <code>ebx.servlet.http.useHttpsSettings</code> was set to <code>true</code>, they could be redirected to the HTTPS page. This behavior is now disabled by default because it could cause issues if EBX® is behind a reverse proxy. To re-enable this behavior, you need to set the new property <code>ebx.servlet.http.authentication.redirectToHttps</code> to <code>true</code>.</p> <p>See <i>TIBCO EBX® main configuration file</i> for more information.</p> |
| Logs to the console are more verbose | All logs, that are not for monitoring or auditing, sent to files are now also sent to the console. |

6.2.0 Known issues

Data Services

- **CP-26823** When the SOAP header input parameter value contains multiple CDATA content it can be truncated.

Workaround: When a text value contains an XML document, the entire contents must be in a CDATA content.

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.

6.2.0 Closed issues

Container Edition

- CP-26805 The docker-buildx version should always be set.
CP-26984 Legacy addon installation in batch mode aborts with the message 'does not exist'.

Core Engine

- CP-4631 [18632] After column reactivation, replica content for On-Demand replication may not be accurate.
This can result in reactivated columns remaining empty after refreshing an OnDemand replication unit.
- CP-22389 Errors can occur during archive export and import.
- CP-25648 `ValueContext.getValue(aPath)` returns null when using an absolute path during the validation of a `TableRefFilter`.
- CP-26312 An error can occur when searching with a value greater than `Integer.MAX_VALUE`.
In some environments, an error is randomly displayed when creating a filter criterion over an integer field with a value greater than `Integer.MAX_VALUE`.
- CP-26427 An EBX instance can initialize indexes even if the folder containing the indexes is locked by another instance.
- CP-26454 The RRO heartbeat thread is stopped too early during EBX shutdown.
As the repository shutdown time can be long if there is several index to commit on disk, the standby node can start before the complete shutdown of the main node. This will corrupt the index.
- CP-26482 It is possible that blocks of records are committed with a transaction id higher than the table's last transaction.
- CP-26484 An incorrect result is returned when applying an empty/is not empty operator on a multi-value foreign key to a hierarchy view.
- CP-26503 It is rare, but possible, that corruption can occur during an index rebuild.
- CP-26539 Validation purge should be more robust.
- CP-26571 The behavior of predicate matches between Lucene optimized and non-optimized filters is inconsistent.
- CP-26590 Stop words should be case insensitive.
- CP-26668 Tables that use the Jaro-Winkler search strategy are not sortable.
- CP-26669 Ensure the repository directory global lock can handle a misconfigured hostname.
- CP-26672 `RtcTableRefUserInputFilter` does not correctly evaluate the `tableRef` programmatic label.
- CP-26676 A `CompileException` occurs on Calcite dynamic code after executing a quick search on an association field.
- CP-26699 The index searcher is not always refreshed to the last committed state.
This can cause the loss of the last index searcher without notice, which can result in queries that return invalid results, such as missing records, and exceptions for deleted records in blocs.
- CP-26709 A `PathAccessException` error is thrown when revalidating a dataset which contains a linked field that targets a different dataset.
- CP-26715 A `PrimaryKeyFilter` inverse with many values can lead to a `Lucene IndexSearcher $TooManyNestedClauses` exception.
- CP-26780 Performance for small, high frequency transactions is degraded.
- CP-26790 After a server restart, the initial version of a reopened dataspace has no datasets.
- CP-26795 An incorrect result is shown in the combo box of a foreign key field (`RtcFacetOnTableReference#getUserInputFilter`) if the label contains another FK label.
- CP-26817 An exception occurs when accessing the history content of a table that contains a FK with a filter referring to a computed value.
- CP-26915 A quick search returns incorrect results when a FK (no label) contains another FK (with label).
- CP-26931 A `NullPointerException` occurs in `LxiUtils.rebuildLuceneIndex()`.
- CP-26936 A deadlock is possible on `LvnTracker` when operating on the index tree structure.

- CP-27172 It is very unlikely, but possible that a `NullPointerException` can occur when sharing an index.
- CP-27180 `Trie.visitrange()` doesn't work with ids that span more than 7 bytes.

Data Modeling

- CP-22251 **Performance can be poor when editing a data model.**
This is observed when editing a data model in the data model assistant, particularly when it contains a significant number of fields.
- CP-26621 **Built-in triggers are imported as custom triggers in the DMA.**
- CP-26809 **A filter predicate set on a foreign key throws an `Invalid XPath` expression error.**

Data Services

- CP-26770 **When the SOAP header value exceeds a common size, the content is truncated.**
- CP-27066 **The operation code is not handled correctly in history mode and is not returned in the response.**

Data Validation

- CP-25217 **A model can still be imported when it contains forbidden characters (e.g. '@') on field names.**
- CP-25576 **When validating a dataset in a child dataspace, the following validation error can occur: "Unavailable record with primary key ...".**
- CP-26350 **An `IllegalStateException` occurs when reading a persisted validation message on a record.**
- CP-26489 **Validation messages related to hidden fields are displayed in record forms.**

Data Workflows

- CP-26697 **The wrong arguments are used for a cached helper class with `CompilationCacheMode.OPTIMIZED_IMPROVED`.**

Hierarchies

- CP-26413 **An error occurs when applying a hierarchy view for a table with multi-value recursive foreign key dimension.**
- CP-26577 **A ragged hierarchy with a recursive join link is not pruned correctly.**

Infrastructure

- CP-26560 **The scheduler task aborts if the repository has a shutdown using failover architecture.**

Scripting

- CP-21458 **Dependencies on computed values are incorrect.**
The current implementation does not define dependencies. As a consequence, indexes can be erroneous if the rule references a field that uses a `ValueFunction`.
- CP-22762 **An EBX script may incorrectly check the integer function field return type.**
- CP-26735 **The primary key type for `deleteRecordByPrimarykey()` cannot be set.**

Team Up

- CP-26954 **A blank page is shown on the table when Teamup is enabled on an existing dataset.**

User Interface

- CP-17286 **A display error occurs while accessing a perspective action on a user's favorited perspective when the scope is set to data.**

| | |
|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CP-23150 | Displaying a table view throws a <code>PreconditionFailure</code> exception. |
| CP-25589 | The 'User logged out' page is shown after a workflow finishes or is cancelled in the user service. |
| CP-25977 | In an empty association, an error occurs in the UI where the text "No records found." overlaps the buttons at the bottom of the screen. |
| CP-26062 | An error occurs when clicking the 'Clear' button and then the 'Open drop-down list' button next to the 'Layout' field on the 'Create a new record' screen on 'Toolbars'. |
| CP-26125 | The 'Access data' service default behaviour was changed and the option is missing. |
| CP-26180 | The wrong perspective is requested in a user service that is nested in a workflow step and perspective. |
| CP-26359 | An error occurs when applying a hierarchy view enabling the 'Prune hierarchy' and 'Allow children count' properties. |
| CP-26455 | In certain cases, the workflow inbox is not shown when a workflow ends. |
| CP-26612 | The default login page encoding is incorrect when the JVM is not set to UTF-8. |
| CP-26678 | Programmatic labels are ignored by the selectors in the search pane. |
| CP-26691 | Forced SSL mode returns HTTP in the login page instead of HTTPS. |
| CP-26741 | Negative session durations should be handled. |
| CP-26761 | The application locking page is not displayed. |
| CP-26766 | The method <code>setFirstCallDisplayHierarchyExpanded()</code> should not be ignored. |
| CP-26807 | It isn't possible to hide the 'Information' option from specific users. |
| CP-26897 | The sidebar does not work on Apple Safari. |
| CP-26911 | Pop-ups displaying records of hierarchical views don't close. |
| CP-26950 | The login page should not redirect to HTTPS even if <code>ebx.servlet.http.useHttpsSettings</code> is set to true. In order to avoid confusion, a new <code>{_}ebx.servlet.http.authentication.redirectToHttps{__}</code> property should be added. |
| CP-26967 | The login page can generate an invalid URL. This can occur if the <code>ebx.servlet.https.authentication.path</code> property has been set and <code>ebx.servlet.http.authentication.path</code> has not been set. |

Users & Permissions

| | |
|----------|------------------------------------------------------------------------|
| CP-26851 | The computation of table level service permission should be optimized. |
|----------|------------------------------------------------------------------------|