TIBCO EBX®
Release Notes

Version 6.2.1
March 2025

Table of contents

6.2 service packs

6.2.1

6.2.1 has been built on March 19, 2025. Overview of this service pack:

Indexing performance improvements	**Optimized index rebuild chain:** An index rebuild now skips steps to shorten rebuild chains, improving the rebuild duration and index structure by reusing base indexes more effectively. This change reduces rebuild times for intermediate snapshots; however, accessing skipped snapshots might take slightly longer, as their indexes must be rebuilt incrementally from the nearest available point.
	Dynamic assessment of incremental rebuild viability: EBX® now estimates the "distance" between table revisions to determine whether to perform an incremental, or a full index rebuild. A full rebuild is initiated when the difference between the base and target index is too large. This adjustment leads to faster rebuilds and more efficient indexing. Accessing a snapshot older than this breaking point can take longer, as a new rebuild chain must be computed from scratch.
	Batch processing for index rebuilds: Index rebuild operations are now batched, grouping deletion and creation actions for greater efficiency.
	Configurable multi-threaded rebuilds: The ebx.index.threads property allows you to control the number of threads used during index rebuilds. Increasing the number of threads can accelerate index rebuilds, especially in contexts with many large indexes after update, migration, or significant data model changes. However, for normal operations, the default setting of '2' threads is recommended.
	For additional details, see *Architecture and design* in the *Performance and tuning* chapter.
Core engine	The memory footprint of the probabilistic filter used to speed up primary key lookups has been reduced by raising its false positive probability. The better accuracy was not worth the additional memory consumption and hash functions to check.
	Some operations have been improved to avoid loading heavy objects that represent datasets when it is not necessary. This can enhance for instance requests on foreign key and inherited fields that need to access to datasets to resolve the target table of the involved relationships.

Query limitations	Performance was optimized when managing all fields in an inherited dataset. See Query limitations for more information.
EBX® Script	It's now possible to create dataset triggers. A dataset trigger defines procedures that are implicitly and automatically executed when some operations are performed on the dataset. See *Dataset trigger* for more information.
Single Sign-On (SSO)	EBX® now supports both SAML 2.0 and OpenID Connect protocols. To enable these features, EBX® should be declared as an application in the Identity Provider and configured depending on the enabled protocol. See *Single Sign-On (SSO)* for more information.
SCIM feature	Built-in support for SCIM (System for Cross-domain Identity Management) was added. The SCIM feature automates the provisioning, management, and synchronization of user identities and roles between an IdP (Identity Provider) and EBX®. See *Enabling the SCIM feature* for more information.
EBX® Container Edition	It's now possible to configure Single Sign-On (SSO), Bearer authentication, the SCIM feature, LDAP built-in directory, and custom directory using environment variables. See *SSO SAML 2.0 connectivity*, *SSO OpenID Connect connectivity*, *Bearer Authentication*, *SCIM feature*, *LDAP connectivity* and *Directory connectivity* for more information.
Data Services	It's now possible to use 'Bearer Authentication Scheme method' with the built-in RESTful services, SOAP operations and custom REST services (built using REST Toolkit). The REST Toolkit *Exception handling* section was updated with the description for a custom exception mapper. See *REST authentication* and *SOAP authentication* for more information.
Supported databases	Support for the following database was added: **Microsoft Azure Database for PostgreSQL** (compatible with PostgreSQL 16 or higher).

	PostgreSQL 17 is now supported and PostgreSQL 12 is no longer supported. The minimum supported version of PostgreSQL is now 13.
	See *Supported databases*, for more information.
User interface validation	A new option has been added to customize the behavior of APV.
	See also *APV should store invalid input in DataContext*
Security - Login page	When a user logs in or out and a custom URL is provided, EBX® appends a `resume` query parameter. It is now possible to disable this behavior by using the following properties:
	• `ebx.security.loginPage.appendResumeUrl`
	• `ebx.security.exitPage.appendResumeUrl`
	• `ebx.security.exitErrorPage.appendResumeUrl`
	See *Security configuration* for more information.
Security - HTML headers	New options have been added in order to include additional JavaScript files or to customize the CSP.
	• `ebx.security.csp.script-src`
	• `ebx.security.csp.style-src`
	• `ebx.security.csp.img-src`
	• `ebx.html.header.javascript`
	See *Security configuration* and *HTML headers* for more information.
History view of a record on the sidebar	It is now possible to display both simple and complex data types, including multi-valued list types.. Additionally, it is possible to filter the history transactions by field.
HTML input	Previously, *CKEditor* was the library used for HTML inputs.
	For security and licensing reasons, it was replaced by a widget based on the open source tool *TipTap*.
	As a consequence, the look and feel has changed. As outlined below, some features were moved or deleted:
	Look & feel
	• Overall style has been reviewed
	• Editor size is not fixed anymore, it now adapts to its content.

Limits can be customized with the following four new settings in the UI administration's *Ergonomics and layout*: **Forms: minimum width of HTML editor**, **Forms: maximum width of HTML editor**, **Forms: minimum height of HTML editor**, and **Forms: maximum height of HTML editor**.

- Editor height cannot exceed the workspace content height.

- Typing assistance has been added, for example if a user types "- foo", it transforms the text into an itemized list.

Toolbar

- Font size menu has been removed in favor of a heading menu in which user can choose between different heading levels.

- A strike-through text style button was added.

- The copy/cut/paste toolbar was removed (keyboard shortcuts and contextual menus are still supported).

- The link toolbar was removed, but links are still supported. This is now transparent for users.

  Links are now created via copy/paste. If a user pastes a link while selecting text, the selection become a link. If the user does not have a text selection, a link is created with its text as the link value.

- A block quote text style button was added.

- Text background and foreground color menus were reviewed.

- The print and preview buttons were removed.

- The maximize button was removed.

- The image menu was moved inline with the editor content; it appears when the user clicks on an image.

- Toolbars can no longer be expanded; it was reduced to a single line.

Background tasks

Configured tasks can now run in the background, which can free up users to continue working in EBX® without having to wait for the task to complete. A centralized user interface (background task sidebar) allows users to view and access these tasks.

See *Background tasks* for more information.

Validation policy for table view

It is now possible to define a validation policy to be used when displaying records in table view. This option can be specified in the default view of the associated table and/or in the global properties of the data model for all tables.

For more information, see :

- ValidationPolicyForTableViewEnum^{API}
- SchemaDefaultTableView^{API}
- SchemaNode.getSchemaDefaultTableView^{API}

Highlight of the quick search results	The quick search results are now highlighted for better visibility.
New software delivery site	You can now download EBX® software from *TIBCO eDelivery*. If you have an existing eDelivery account, your existing credentials can be used to login to the new software delivery site.
Known issues	See <u>6.2.1 Known issues</u> section.
Closed issues	See <u>6.2.1 Closed issues</u> section.

Java API

The Java API was improved on the following topics:

- WidgetDisplayContext

  New method **WidgetDisplayContext.isInherited()** WidgetDisplayContext.isInherited^{API}.

- PerspectiveTheme

  New methods **getButtonsHighlightTextColor()** PerspectiveTheme.getButtonsHighlightTextColor^{API}, **getFormHTMLEditorsMaxHeight()** PerspectiveTheme.getFormHTMLEditorsMaxHeight^{API}, **getFormHTMLEditorsMaxWidth()** PerspectiveTheme.getFormHTMLEditorsMaxWidth^{API}, **getFormHTMLEditorsMinHeight()** PerspectiveTheme.getFormHTMLEditorsMinHeight^{API}, **getFormHTMLEditorsMinWidth()** PerspectiveTheme.getFormHTMLEditorsMinWidth^{API}, **getPerspectiveNavigationPaneSelectionTextColor()** PerspectiveTheme.getPerspectiveNavigationPaneSelectionTextColor^{API} and **getWorkflowBadgeTextAndBorderColor()** PerspectiveTheme.getWorkflowBadgeTextAndBorderColor^{API}.

- ValidationPolicyForTableViewEnum

  New enumeration **ValidationPolicyForTableViewEnum** ValidationPolicyForTableViewEnum^{API}.

- SchemaDefaultTableView

  New interface **SchemaDefaultTableView** SchemaDefaultTableView^{API}.

- SchemaNode

  New method **SchemaNode.getSchemaDefaultTableView()** SchemaNode.getSchemaDefaultTableView^{API}.

- DifferenceHelper

  New methods **compareDatasetByDate(Adaptation,boolean,java.util.Date,java.util.Date)** DifferenceHelper.compareDatasetByDate^{API},

compareDatasetByTransactionID(Adaptation,boolean,long,long) `DifferenceHelper.`
`compareDatasetByTransactionID`^{API},
compareDataspaceByDate(AdaptationHome,boolean,java.util.Date,java.util.Date)
`DifferenceHelper.compareDataspaceByDate`^{API},
compareDataspaceByTransactionID(AdaptationHome,boolean,long,long)
`DifferenceHelper.compareDataspaceByTransactionID`^{API},
compareRecordByDate(Adaptation,boolean,java.util.Date,java.util.Date)
`DifferenceHelper.compareRecordByDate`^{API},
compareRecordByTransactionID(Adaptation,boolean,long,long) `DifferenceHelper.`
`compareRecordByTransactionID`^{API},
compareTableByDate(AdaptationTable,boolean,java.util.Date,java.util.Date)
`DifferenceHelper.compareTableByDate`^{API}   and
compareTableByTxId(AdaptationTable,boolean,long,long) `DifferenceHelper.`
`compareTableByTxId`^{API}.

- `MergeSpecForInstanceContents`

  New constructor **MergeSpecForInstanceContents(boolean)** `MergeSpecForInstanceContents.`
  `%3Cinit%3E`^{API}.

- `UserTask.WorkItem`

  New   method   **UserTask.WorkItem.isCommentEnabled()**   `UserTask.WorkItem.`
  `isCommentEnabled`^{API}.

- `WorkflowEngine`

  New   method
  WorkflowEngine.getProcessInstance...History(PublishedProcessKey,String,String)
  `WorkflowEngine.getProcessInstanceKeysForDataContextVariableIncludingHistory`^{API}.

*Note: **6.2.0** has been built on June 21, 2024.*

Workflows

Notable points done for workflows:

Workflow views performance	• Improvement of performance for specific columns in Data workflows views. Specific columns should be sortable and filtering on them is now faster. But a tradeoff was made: Built-in workflow variables are no longer supported, only data context variables are supported. The expression is now evaluated each time the workflow data context is updated, not on display but the result should be seemless for the user with the previous behavior.
	• New property added to replace Label (dwfLabel) for Workflow related tables (Active/Completed workflows and in Workflow administration).
	New property added to replace Current publication name (dwfPublicationLabel) for publication in publication related tables (Publication and in Workflow administration).
	The new properties take the current previous property localized label while the old are added deprecated. The new properties are now correctly filterable but still not sortable. The filtering is now faster.
	Some other properties are impacted by those changes: Data workflows in Inbox/Work item view, Publication label in Inbox,Work item, Active/Completed Workflows and Workflow Parent in Active/Completed Workflows views. They are mapped on the new properties and are faster to filter but sorting is still not resolved, for the time being the sort is done by primary key by default.
	When using a filter like contains that search a part of label you might try doing a AND filter with the second filter being a search filter with the same value. It might return a result faster.
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.

Performance

A prefilter can now be defined for an AdaptationFilter[API] or TableRefFilter[API] 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.

REST Toolkit APIs

The *REST Toolkit* now supports *Jakarta RESTful Web Services 3.1*.

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` relation ship 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 *LDAP directory* and *Configuring built-in LDAP directory*, 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 `PrefilteredAdaptationFilter`[API] has been added for adding a prefilter to an `AdaptationFilter`[API].

A `TableRefFilter`[API] can define an XPath prefilter during the setup of the table reference filter. See `TableRefFilterContext.setPrefilter`[API].

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 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. **See also** *Overview of the Script IDE*
ebx-hub.war	This is the WAR for the new Hub web application. **See also** *Home page*

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.
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`	A nested procedure now cannot use the `ProcedureContext` 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.
	Attempts to operate on the `ProcedureContext` of the parent procedure will now result in an `IllegalStateException`.
Data validation	Internal validation framework has been enhanced to minimize the number of snapshots created to validate data. It now uses as-of transaction time to compute the delta between records, tables, datasets or dataspaces. This simplifies the validation process by avoiding to create snapshots for functional dataspaces involved in this process. This optimization also improves purge execution time by considerably reducing the number of snapshots that will be deleted.
	As a consequence, theses changes will have the following impacts:
	• The validation reports created by EBX® versions prior to 6.2.1 are not compatible and will be reset. The stale data for the reset validation reports will be removed during the next purge.
	• The validation reports are no longer copied when a dataset is duplicated. It will therefore be fully validated for the first time.
	• An error can be raised if a constraint that is being checked during the validation of a dataset or a table execute a procedure in the same dataspace.
Optimization of matches predicate	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 ^ (beginning of line) or $ (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.
Association with inconsistent definition	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.
As a consequence of this change an exception is no longer raised in the existing API that compute the result of the association.
See `AssociationLink.getAssociationResult`[API] and `AssociationLink.getTargetTable`[API]
For more information, see *Association*.
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	Since EBX® 6.1.3 HF-1, when users accessed the login screen using HTTP and the property `ebx.servlet.http.useHttpsSettings` was set to `true`, 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 `ebx.servlet.http.authentication.redirectToHttps` to `true`.
See *TIBCO EBX® main configuration file* for more information.
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.

6.2.1 Known issues

User Interface

- **CP-24388** Tile 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 tile view cards.

6.2.1 Closed issues

Container Edition

CP-26866	**The `EBX_SMTP_CONNECTION_TIMEOUT` and `EBX_SMTP_WRITE_TIMEOUT` environment variables are ignored.** When the SMTP timeouts properties are configured with new value, the value is only applied for the `EBX_SMTP_TIMEOUT` parameter.
CP-27492	**The Container Edition installer can abort with `docker: invalid reference format` on Windows platform.** When the installation folder contains a space character, the installer aborts.
CP-27850	**When the resources cache size is set to a value that is too small, a warning is logged at startup.**

Core Engine

CP-17505	**For inherited datasets, a search operation is not pushed into Lucene.**
CP-26118	**A "read only" merge is not applied to all table revisions at child dataspace creation.** `SetShared()` is only called on a cached table revision when a child dataspace is created. A "read only" index merge is never applied to other revisions.
CP-26606	**The date formatter used by computation rules is not thread safe.**
CP-26689	**A case insensitive XPath match can produce incorrect results.**
CP-26928	**A `SortJoinTransposeRule` is continuously triggered due to incomplete Project collations, which leads to a `StackOverflowError`.**
CP-26932	**A `NullPointerDxception` occurs in `LxiUtils.rebuildLuceneIndex().`**
CP-26993	**There is an inconsistency in the free slot state of a block.** In the case where blocks are full, but have deleted records, the `BlkFreeSlotState#hasFreeSlot` method can return `true` while the `BlkReadWrite#firstFreeSlot` method returns `BlkBloc.NO_FREE_SLOT`.
CP-26996	**The `CyclicDependencyChecker` should be improved.**
CP-27111	**Teamup's dataset is not deleted after disabling the related extension in the data model.** Teamup's technical dataset is not removed when deleting a dataset after having disabled the extension in the underlying data model. As a consequence teamup's dataset remains in a state of severe error in the dataspace where it has been used and it is not possible to perform some operations from this dataspace (such as merging it to its parent).
CP-27137	**Updates set to the `BeforeCreateInstanceContext` dataset trigger are not performed.**
CP-27206	**An incorrect Calcite hierarchy plan in `LpRecursiveBinaryLinkTargets` can throw an exception.**
CP-27215	**The method `Adaptation.getOccurrencePrimaryKey()` does not return `null` when called on a dataset.**
CP-27227	**A column in a historized table cannot be purged after disabling and purging the history table.** An error occurs when trying to purge a disabled column that has been previously renamed in a historized table. This issue occurs on tables that are replicated and historized at the same time. It happens after disabling the history, then purging the history table, and re-enabling the history. In this case the purge of the column fails because the renamed column does not exist in the history table that was created when re-enabling the history.
CP-27235	**Dataspaces reopened after a server restart are not properly initialized.** This issue can prevent existing datasets from loading, which leads to them being recreated. This can cause a corrupted repository error on startup because the same dataset reference exists multiple times.
CP-27266	**The `IndexSearchers` used by active concurrent transactions were not properly retained.** This can cause the following kinds of errors: Index already closed, Cannot find index state at transaction, and Cannot fetch record at position.
CP-27295	**A `MinShouldMatch` causes records to be wrongly considered as not similar.**
CP-27317	**Desynchronization between index and persistence can lead to a `MissingRecordException`.**
CP-27337	**The index is opened on an outdated transaction after a rollback or truncate operation.**

CP-27389	**An error occurs when renaming a column in the auxiliary table related to an aggregated list.**	

This issue happens when changing the data type of simple list elements, purging the previous column, and renaming the new one with the previous name.

CP-27416	**Archive imports do not properly set the last transaction.**

Any discrepancies can lead to errors when rebuilding indexes, or performing accesses or comparisons by transaction id.

CP-27431	**A purge of shared blocs can saturate the DBMS transaction log.**

The purge might lead to deleting a large number of blocs from the database, overwhelming the database's transaction log. Intermediate commits must be introduced to make the process lighter for the underlying DBMS.

CP-27465	**Computing the proximity between table revisions is too expensive in terms of computational resources.**
CP-27485	**Ensure all dataspaces to reopen are loaded in the cache before processing.**
CP-27502	**Optimization is not always applied to the index during an incremental rebuild.**

In the index rebuild chain, some revisions are seen as unshared even if they had children. This is usually the case for snapshots. In this case, the optimized merge policy for a read-only index is not applied. For a long rebuild chain this will create sub optimal indexes with too many leaves.

CP-27526	**A `NullPointerException` occurs when starting a server and registering modules.**
CP-27529	**The wrong Logging context is picked up on a module's logging configuration update.**

The thread caller logging context is used instead of the module's.

CP-27547	**Performance issues can occur when displaying the list of scheduled tasks if the directory contains many users.**
CP-27563	**An association cannot be created that is based on a linked field that inherits from a foreign key field.**
CP-27580	**Index rebuild during a Procedure must not commit pending changes.**

When an index rebuild happens during a Procedure, the pending changes of the running transactions are committed in the index with the last committed transaction id of the table. This can lead to discrepancies when:

* The procedure that triggered the rebuild is eventually canceled; the canceled changes remain persisted in the index.

* A concurrent access to the table before the Procedure is committed will see the pending changes in the index even though they are not accessible in the cache.

CP-27594	**An `IndexSearcher` might be lost when it is refreshed in a Procedure with nested sub-Procedures.**

This can occur after committing a nested sub procedure.

CP-27630	**The underlying `IOError` is not propagated when rebuilding the index.**
CP-27650	**A dataset header cannot be accessed if it was imported with the wrong transaction id.**
CP-27669	**An `EnumerableSortedNestedLoopRight` error occurs without collation.**
CP-27682	**A performance issue occurs when accessing a custom views dataset in the administration area.**
CP-27702	**In rare cases index corruption can occur after EBX® stops abruptly.**

This can also occur with a dataspace chain that has very few modifications in the parent index. The internal file sequence of a child index without any commit relative to it's parent can be wrong.

CP-27762	**The Javadoc for `ProcedureContext` is malformed.**
CP-27897	**The background deleter can flood `persistence.log` on the Windows platform**
CP-27985	**Child dataspace creation retains an index searcher longer than is necessary.**

The index searcher for the last transaction of every loaded table of the parent dataspace is retained. This consumes more memory resources than necessary. With lower amounts of free memory, this can lead to an out of memory issue, or leaked file resources.

CP-27992	**Transaction rollback retains an index searcher for the last transaction longer than necessary.**

The index searcher for the last transaction of every index is retained after rollback. This consumes more memory resources than necessary. With lower amounts of free memory, this can lead to an out of memory issue, or leaked file resources.

CP-27993	**An `IndexSearcher` can lead to a memory leak.**

In some cases, `IndexSearchers` are not detached from `LxiStateWithWriter.lastRefreshedInProcedure`. This prevents them and the underlying index from being closed, leading to memory and mapped file resource leak.

Data Modeling

CP-26898	**Performance is slow when accessing the service to import a data model from a module.**
	This can occur if a lot data models are present in the modules registered in the application server.
CP-27643	**The system should avoid generating a temporary data model when validating toolbars in the DMA.**
	The data model is validated prior to its publication. A performance issue can occur when validating a data model that uses a lot of toolbars. In this case, a temporary XSD is generated for each check.
CP-27948	**The schema generated when publishing the data model should not be cached if it contains errors or it doesn't include Add-ons.**

Data Services

CP-24734	**The built-in `prepareForCreation` operation aborts with a 500 error status code under certain conditions.**
	This can occur when a primary key field type is `xs:date`, `xs:time`, or `xs:dateTime`.
CP-27372	**When using the REST API's built-in OpenAPI to perform a 'select' operation in compact mode, the association fields are not included in the response.**
CP-27420	**REST Toolkit exception handling can abort and return an HTTP 500 error.**
	If the Java REST service uses built-in exception handling and creates a method that does not declare the `@Produces` Java annotation with an "application/json" media type value, then an error from the service creates an HTTP 500 with the following message: 'No message body writer has been found for class'
CP-27603	**REST built-in 'select' operations do not return all record levels on business objects.**
CP-27611	**There is no JSON body response with the built-in 'select' operation when the query parameter `includeMetadata` includes an incorrect value.**
	This occurs even if the `Accept` HTTP header is defined with: `*/*` or `application/json`.
CP-27617	**REST built-in 'select' operations on inheritance can have smaller responses with `includeOcculting` and `includeMetadata`.**

Data Workflows

CP-26999	**Random duplicate entry issues occur when publishing a workflow model.**
	The index recovery process on a table with record versioning disabled may result in an inconsistent index and discrepancies between the index and the persistence. Recovery that normally happens after a forced shutdown or a procedure rollback must be disabled. As a consequence, the index will be rebuilt. When the rebuilding operation runs concurrently with the execution of a procedure, the running procedure is canceled. These limitations remain acceptable since disabling record versioning is an internal property that is only applied on a few technical tables.
CP-27204	**A staging error occurs when the name of a workflow is empty.**
	The `StagingException: Function identifier without id is not supported` occurs when accessing the staging widget after creating a workflow with an empty name.
CP-27399	**The format for `processInstanceKey` returned by `getProcessInstanceKeysForDataContextVariable` is incorrect. method.**
CP-27864	**Workflow models created by EBX® version 5.9 using script tasks might have an incorrect mode type.**

Documentation

CP-27325	**The documentation is missing some SOAP operations and REST comparisons.**
	The missing information includes select business objects, health operations, and generating custom service contracts.
CP-27340	**The environment variable names are inconsistent for the Metadata database.**
CP-27918	**Update the REST Toolkit supported version in the documentation.**

Hierarchies

CP-20093	**The hierarchy-specific mode 'Display non-matching children' is not implemented.**
CP-26945	**A blank screen is displayed while selecting a level configuration panel.**
	In some rare multi-dataspace, multi-dataset configuration A blank screen is displayed due to a `NullPointerException` in the method `UISimpleFilterHelper.isSearchable()`.

CP-26962	**Hierarchy dimension introspection can cause performance issues.**
CP-27261	**A `ClassCastException` occurs when closing a hierarchy record panel.**
	When the hierarchy dimension consists of two consecutive recursive levels, an exception can occur when closing a record panel or using the Find action.
CP-27362	**Hierarchy performance issues can occur on child dataspaces.**
	If a workflow creates a temporary dataspace when displaying a hierarchy, the Calcite request cache needs to be rebuilt each time. This can cause the performance issues.
CP-27489	**A `PrimaryKeyException` occurs when deleting a hierarchy member.**
	In a ragged hierarchy, a `PrimaryKeyException` can occur after deleting a recursive member.
CP-27608	**Children of searched nodes should not be expanded.**
	When the corresponding option is set, children of searched nodes should not be expanded.

Import/Export

CP-27567	**Reserved characters should be percent-encoded in the HTTP Header Content Disposition.**
CP-27625	**The hierarchy member label is encoded when exported to a CSV file.**
	When a CSV is exported from a hierarchy member service, the character '&' is encoded, although it should not be.

Infrastructure

CP-27276	**The module log level is set to DEBUG when using the module log configuration syntax from previous versions.**
CP-27277	**The custom module log threshold is not used when an `ebxFile` appender is specified.**
	The INFO threshold is used instead.
CP-27571	**The `DirectoryHandler` Java API must return immutable collections.**

Metadata

CP-27030	**Metadata services are not registering correctly in some Unix environments.**
CP-27148	**Users are redirected to a blank window when creating data sources while the metadata harvesting service for instances is running.**

Scripting

CP-27297	**The built-in Script IDE does not support data models that have multiple versions.**
CP-27374	**A function or trigger script that declares an alias on its own schema might not compile.**
CP-27473	**The built-in Script IDE does not support data models that have multiple versions.**
CP-27694	**The script IDE throws an error when the TIBCO EBX® Match and Merge Add-on is activated.**
CP-27916	**The EBX® script documentation for methods that lookup dataset fields should be updated.**

Staging

CP-26893	**An error occurs when staging a data model definition from a Windows server to a Linux server.**
CP-27369	**The staging widget is blocked due to an incorrect view label calculation.**

User Interface

CP-22144	**A browser confirmation dialog to leave the page can sometimes display after modifying and publishing a data model.**
CP-23196	**The tooltips of some fields aren't properly shown in the task scheduler's screens.**
CP-25607	**The Information sidebar service does not work for old datasets.**
CP-26808	**Ajax prevalidation should not block context updates.**
CP-26839	**When a field has a custom widget, its default value cannot be cleared.**
CP-27092	**An error occurs when creating a custom view filter using a linked field.**
	In the configuration screen of a custom view, when creating a filter using a linked field linking to an enumeration, the 'Direct selection' field does not work.

CP-27232	**In a user service, an error is thrown when deleting a record in an association from another dataset.**
CP-27293	**The permission service can throw a `NullPointerException`.**
CP-27371	**Sidebar history fails when the model includes a folder with no terminal nodes.**
CP-27456	**Navigation is not handled correctly when cancelling a dataspace export.**
CP-27482	**Some UI pages might display an incomplete product version.**
CP-27497	**White spaces in record form fields should be preserved.**
	In the form view, text fields which are displayed with a custom widget and are configured to be read only should display the persisted value and preserve white spaces.
CP-27554	**In tile views, the enumerated values are not displayed correctly.**
CP-27562	**The session is not closed when clicking on the 'X' button.**
CP-27662	**The search criteria tree structure might get ignored.**
CP-27674	**In certain cases, an incorrect error message is displayed for users.**
CP-27818	**The memory consumption of QSAObjects should be reduced.**
CP-27859	**The HTML input should be changed to a JavaScript bundle extension.**
CP-27892	**An error that occurs when loading a theme prevents the UI from displaying.**

Users & Permissions

CP-27219	**Setting permissions to hide an association tab in child dataspace does not work as expected.**

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 in 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 `NullPointerDxception` 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 `PreconditionFailure` 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 `setFirstCallDisplayHierarchyExpanded()` 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 `ebx.servlet.http.useHttpsSettings` is set to true.**

In order to avoid confusion, a new {_}ebx.servlet.http.authentication.redirectToHttps{_} property should be added.

CP-26967 — **The login page can generate an invalid URL.**

This can occurs if the `ebx.servlet.https.authentication.path` property has been set and `ebx.servlet.http.authentication.path` has not been set.

Users & Permissions

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