



TIBCO EBX®

Release Notes

Version 6.2.3
April 2026

Table of contents

6.2 service packs.....	5
6.2.3.....	5
6.2.2.....	11
6.2.1.....	16
Workflows.....	22
Performance.....	24
Data validation.....	24
Application servers.....	24
Data services.....	25
REST Toolkit APIs.....	25
Built-in RESTful services.....	25
Built-in LDAP directory.....	25
Improving the performance with prefilters.....	25
EBX® Container Edition.....	26
EBX® Script IDE.....	26
Home page.....	26
Metadata management.....	26
History view of a record on the sidebar.....	27
Product migration.....	27
Limitations for this release.....	27
Hierarchy limitations.....	27
Query limitations.....	27
Value-labeling limitations.....	28
Data validation limitations.....	28
Version upgrade.....	28
New Web Application Resource (WAR) files.....	28
Compilation and build environments.....	29
Changes on supported environments.....	29
Upgrading from an existing EBX® 6 environment.....	29
Backward compatibility.....	31
6.2.3 Known issues.....	33
User Interface.....	33
6.2.3 Closed issues.....	34
Assistants.....	34
Core Engine.....	34
Data Modeling.....	35
Data Services.....	35
Data Validation.....	35
Data Workflows.....	36
Distributed Data Delivery.....	36
Documentation.....	36
Hierarchies.....	36
Infrastructure.....	36
Mame.....	36
User Interface.....	36
Users & Permissions.....	36
6.2.2 Closed issues.....	38
Container Edition.....	38
Core Engine.....	38
Data Modeling.....	38
Data Services.....	39
Data Validation.....	39

Data Workflows.....	39
Documentation.....	39
Hierarchies.....	39
Infrastructure.....	39
Scripting.....	39
User Interface.....	39
Users & Permissions.....	40
6.2.1 Closed issues.....	41
Container Edition.....	41
Core Engine.....	41
Data Modeling.....	43
Data Services.....	43
Data Workflows.....	43
Documentation.....	43
Hierarchies.....	43
Import/Export.....	44
Infrastructure.....	44
Metadata.....	44
Scripting.....	44
Staging.....	44
User Interface.....	44
Users & Permissions.....	45
6.2.0 Closed issues.....	46
Container Edition.....	46
Core Engine.....	46
Data Modeling.....	47
Data Services.....	47
Data Validation.....	47
Data Workflows.....	47
Hierarchies.....	47
Infrastructure.....	47
Scripting.....	47
Team Up.....	47
User Interface.....	47
Users & Permissions.....	48

6.2 service packs

6.2.3

6.2.3 has been built on April 29, 2026. Overview of this service pack:

Java Runtime Environment	JRE 25 LTS is now supported. We recommend upgrading your JRE to this latest supported LTS version. <i>See Supported application servers for more information.</i>
Servlet and JSP container	EBX® now requires Jakarta EE 10 and above. This change requires an upgrade for the following third-party libraries: Servlet to 6.0 or greater, JSP to 3.1 or greater, JMS to 3.1 or greater, and Mail to 2.1 or greater. See the links below for more information on specific component requirements. <i>See Supported application servers, Java Message Service (JMS), Installation notes and WebLogic AS support for more information.</i>
New user interface (UI) preview	This release includes a preview of the new EBX® user interface, offering an updated navigation experience and new interaction patterns. The preview is optional, administrator-enabled, and operates alongside the current UI. Additional capabilities will be introduced in future updates. <i>See the Administration Guide and User Guide for more information.</i> Note Use of the new UI is optional. It can be evaluated alongside the existing UI, which remains the recommended interface for production use.
Backward compatibility – Repository API behavior	Starting in this release, you must <i>not</i> call the method Repository.deleteHome(AdaptationHome, Session) <code>Repository.deleteHome^{API}</code> from a procedure <code>Procedure^{API}</code>. If you do, the method throws an <code>IllegalStateException</code>. This change prevents transactional inconsistencies and helps ensure data integrity. Previously, this call could lead to internal errors or leave the repository in an inconsistent state when a nested procedure committed changes before its parent. Attention

This change might impact any custom code that attempts to delete a dataspace (branch or snapshot) from within a procedure. If this is the case with your code, you must move the deletion outside of the procedure context.

Workflow changes

History size reduction

- Data context variables are now *only* recorded with a new history entry if they are updated. This change can lead to a significant decrease of the workflow history size over time.

The total amount of reduction depends entirely on the workflow models and the number of variables updated in each step during their execution. For example: a model of 20 steps with 20 variables defined that updates on average 2 variables per step will reduce the number of stored records by at least 90%.

Note

This change applies only to new workflow execution. Existing workflow history will remain as is.

New built-in variable in Message templates

- Introducing built-in variable 'workflow.workItem.key'. It will be resolved to the work item key in mails.

Workflow user task configuration

You can now set a preference for allocation of the current task to the user who most recently completed an earlier task in the workflow. If the user who completed the earlier task belongs to any profile the current task is offered to, the system allocates (assigns) the task to that user. If that user does not belong to one of those profiles, the system offers the task to users in the eligible profiles as usual. See *Setting a user allocation preference* for configuration information.

Note

This feature is not available for legacy user tasks.

EBX® Script

The language now:

- Allows you to create scripted REST services. These services define a REST endpoint in EBX® and its related script. You create and maintain these services in the Script IDE. After compiling and publishing the

service from the IDE, it becomes available to REST clients.

- Supports validation constraints. Use constraints to ensure data consistency within a data model.
- Allows you to create customized messages with user-defined severity levels. Add the messages to the validation report using EBX® Script.
- Enables creation, modification, and publishing of localized messages from bundles using the EBX® Script IDE.

See also

Scripted REST service

Creating scripted REST services tutorial

Constraint validation

Predefined scheduled task to launch workflow

The new SchedulerWorkflowLaunch task schedules the launch of workflows. The task requires input of a workflow launcher to determine which workflow publication is launched.

Built-in LDAP

The caching system has been improved to support high availability. A new property, [request].cache.refreshInSeconds allows cache entries to be refreshed before they expire.

See *LDAP directory* for more information.

EBX® Metadata Management application

- Beginning with EBX® version 6.2.0 the Metadata Management application has been available when using EBX® Container Edition. You can now install the EBX® Metadata Management application on a traditional (on-premises) deployment.
- The EBX® Metadata Classifier database now supports Microsoft SQL Server and Oracle databases.

See *Installation notes for EBX® Container Edition* and *Building the image* for more information.

- The metadata agent script has been revised following an upgrade to the Spring Framework version. Ensure you use the most recent agent script to connect to Metadata.

See *Installing the TIBCO® Metadata Agent* for metadata agent installation instructions.

EBX® Container Edition

The JDK has been upgraded to LTS version 25.

It is now possible to define the EBX® run mode using an environment variable.

The prefix of the environment variable names for metadata has been updated from `METADATA_` to `EBX_METADATA_CLASSIFIER_`. The old variable names are still supported for backward compatibility but will be removed in a future release.

The customization of the Docker image for metadata classifier, when using databases other than SQLite, has been simplified.

To configure a JNDI mail session for EBX®, you'll need to customize your Docker image.

See *Overview*, *EBX® run mode*, *Metadata classifier database connectivity*, *Adding a driver for the Metadata Classifier database* and *Configuring the mail in JNDI* for more information.

Storage and compression

This release introduces a complete rework of the `G_BLK` format and compression to reduce storage requirements.

- The new format supports partial content sharing with another block (another `G_BLK` table row). This reduces data duplication while preserving isolation across dataspace.
- Compression can now use dictionaries, which improves efficiency. The improvements are especially noticeable on smaller blocks that are more common with the new sharing mechanism.

The storage savings that result from these updates are expected to range from:

- 10% to 25% for freshly imported data.
- Up to 90% for heavily modified data in dedicated dataspace.

You might experience a performance impact as compression speed can now be significantly slower. However, this is not an issue in most cases as:

- There is less data to compress (due to sharing).
- Compression time remains negligible as compared to the full processing pipeline, except for massive initial imports.

If compression performance is an issue, use the following parameters in `ebx.properties` to adjust it:

- `ebx.compression.dictSize` controls the dictionary size and defaults to 16 KB. The recommended range is between 1 KB to 64 KB. However, you should be aware of the following tradeoff that occurs with a larger dictionary: The compression ratio (Logarithmic

progression) is better, but compression is slower (linear increase).

Note

You must set this value ahead of time. After EBX® creates a dictionary for a given table definitions, it does not create another one. The new value only applies to tables that don't have a dictionary, or to new table definitions.

- `ebx.compression.level` controls the compression level. Possible values are FAST, BALANCED, and HIGH (default). Lower the level to speed up compression, but expect a lower compression ratio. This setting immediately applies to the next compression operation. This setting has no noticeable impact on decompression.

For compression performance Java 17 remains fully supported, but running with Java 21 or Java 25 is strongly recommended. The fastest native compression implementation is only available with:

- Java 21 or later, and
- Linux or Windows, and
- x86_64 CPU architecture.

In other environments, compression falls back to a full Java implementation. The fallback is slower and cannot create optimized dictionaries. Both implementations are fully compatible, and you can switch between them.

Developer impact: The optimized compression implementation now uses direct/native memory. This makes it more important to explicitly close `RequestResult` and `QueryResult` objects as soon as possible so that memory (index and database resources) is returned promptly.

Public API behavior change: `ReverseRelationshipResult.getResult()` previously returned the same `RequestResult` instance. It now always returns a new `RequestResult` that you must close after each use.

Installation and deployment artifacts: As of this release the `ebx-1z4.jar` library is no longer needed and has been removed.

Query highlights

Optimizations have been made to the calculation of highlights during searches on linked fields.

Memory consumption has been significantly reduced, especially for large secondary tables. Previously, each search result retained a reference to the full query context, preventing memory cleanup. This could have a major impact when browsing a large number of results pages, as each underlying result was held in memory.

For volumes of several million records, several gigabytes of retained heap could occur, leading to garbage collector blockage.

Key improvements:

- Highlights computation is now deferred using a lightweight structure instead of retaining the full query context
- After computation, references are released for immediate garbage collection
- Retained heap reduction, in some cases, by a factor of 50 to 100x (e.g., from >10 GB to ~100 MB). This depends on each situation.

Supported database	PostgreSQL 18 is now supported and PostgreSQL 13 is no longer supported. The minimum supported version is now PostgreSQL 14. Oracle 21c (Innovative Release) is no longer supported; the required version for Oracle is now 19c or 23c. SQL Server 2014 is no longer supported. The minimum supported version is now SQL Server 2016. See <i>Supported databases</i> for more information.
Session timeout redirect property	A new property allows you to declare where to redirect users after their session has timed out. See <i>Security configuration</i> for more information.
SSO - Data Services	You can now configure which JSON Web Token (JWT) claims are used to identify the authenticated user. For additional details, see <i>OpenID Connect</i> and <i>Security</i>.
Known issues	See 6.2.3 Known issues section.
Closed issues	See 6.2.3 Closed issues section.

6.2.2

6.2.2 has been built on September 2, 2025. Overview of this service pack:

New feature: AI Assistant

A new AI Assistant feature has been introduced to support EBX® users using Generative AI. This capability leverages Large Language Models (LLMs) to help users by generating suggestions based on their input and record data. This is especially helpful for tasks such as generating text, summarizing content, translating values and standardizing formats. To use the AI Assistant, administrators must connect EBX® to an external LLM provider and define assistance functions. Each function includes a prompt, which instructs the LLM on how to generate meaningful results for one or more fields within a record.

For more information, see the following documentation resources:

- *AI Assistant* includes instructions for administrators on configuring an external LLM connection and defining assistance functions.
 - *Editing data* provides directions on using the AI Assistant.
-

EBX® Script

The EBX® Script Language now supports calling external REST services from a script.

Role mapper scripts provide a flexible approach to mapping external roles from identity providers to internal EBX® roles.

For more information, see the following documentation resources:

- *External services configuration* describes the configuration properties and keystore files used to encrypt data.
 - *REST service* includes instructions and examples of calling REST services in a script.
 - *Role mapper*.
-

Synchronized directory

A new synchronized directory has been introduced to handle external users and roles in Single Sign-On (SSO) configurations. To enable integration with a variety of external user directories, the following built-in implementations are provided: LDAP, SCIM and OnLogin. See *Synchronized directory* for more information.

Backward compatibility: For LDAP and SCIM, Manual migration is required to import all roles from an external directory and keep using the same permissions.

See *Migration* for more information.

License keys

The license key manager has been reinstated. Your license key is now aligned with your EBX® instance contract type: subscription, perpetual, or trial. Deployed add-ons are now registered and enabled at system startup.

See *Setting an EBX® license key properties file*, *EBX® Container Edition configuration* and `TestHelper.prepareEnvironmentAPI` for more information.

Data validation

Optimized validation: the validation of the *pattern* facet has been optimized. It now benefits of the use of Lucene indexes and then can be used on large tables. That is, if a field defines this facet then the validation process will no longer iterate over all the records to check if the value of the associated field complies with the specified pattern. See *data validation* for more information about the validation process.

Disable constraint pre-validation: a new method was added to prevent, in the user interface, the constraint from being pre-validated dynamically when the related field is updated at user inputs. This can be useful to improve performance or to avoid unnecessary validation errors during form updates. However, the constraint will still be validated when the form is submitted or when loading the record in the view.

See `ConstraintContext.disableAjaxPrevalidationAPI` and `ConstraintContextOnTable.disableAjaxPrevalidationAPI` for more information.

Workflow changes

New changes on workflow views, [see details below](#).

Java API to retrieve table views

It is now possible to get the table views independently of a UI context. It can be useful for retrieving all views related to a table and for a specific user.

See `UICustomViewContextFactory.newCustomViewContextAPI` for more information.

Java API to retrieve and publish data models

A new Java API has been added to retrieve and publish data models defined in the Data Model Assistant.

See `DataModelingManagerAPI` and `DMADataModelAPI` for more information.

Data Services	New operations have been added to support the publication of data models defined in the Data Model Assistant. See <i>Data model operations</i> for more information.
Tomcat 11 support	EBX® version 6.2.0 and later is not compatible with Tomcat 11.
Performance improvements and resource optimizations	• By reducing intermediate object allocation, CSV import operation speed was increased (throughput +~10%). • Improvements to 'onDemand' replication unit refresh and initial synchronization of the 'onCommit' replication unit result in more efficient resource usage and less intermediate objects and round trips with the database: • Diffs are not computed but inserts are directly applied if the replica can be assumed empty. • Diffs are applied on-demand to the database instead of being stored in memory structures before application. • New index pre-loading configuration option: You can specify when to automatically rebuild dataspace indexes at startup as an asynchronous process. Boot isn't blocked but missing indexes are prepared so they can be ready before first access. See <i>Advanced index configuration</i> for more information. • Index size on disk: • An index rebase and more aggressive merge policy improves index efficiency and shortens the index chain. This makes initialization faster and reduces memory consumption. • Indexes for table revisions that are no longer attached to a dataspace are purged. This reduces the number of indexes, their size, initialization time, and object count. • IndexWriter and IndexReader lifecycle management was improved to ensure that files which are no longer usable are deleted. Note that by default Lucene considers un-deleted files as non-issues because they can get deleted by an IndexWriter the next time it is open. However, with EBX®, some indexes become read-only. Before this improvement, an IndexWriter was not re-opened and the files were never deleted.

- Elias-Fano encoding is now used to store live docs. This is used: on-disk as soon as it can save storage and in-memory when it can save at least 50% of the memory. There is a possibility that some queries can be slowed.
- To save memory, common strings are deduplicated and a more efficient implementation for small, long-live, and nearly immutable collections was applied.
- Generalization of index pooling on Java heap data structures: This is the equivalent to the 6.2.1 improvement to pool the primary key bloom filter on index leaves, which is applied to all index file types. The indexes on Java heap data structures are much smaller than those on the bloom filter. However, the saved memory can keep more indexes open, and increase how much they are shared (or duplicated if pooling is not applied). This update can also improve initialization time since files do not require a checksum each time they are accessed. Additionally, it can help reduce the load on other resources such as OS memory mapped file descriptors.
- Lucene was upgraded from 9.4 to version 9.12.1. This upgrade, and the others mentioned above, require a full index rebuild. However, it will be faster than in previous versions due to index rebuild improvements and the new index preloading service. Lucene 9.12 is currently slated to be the last Lucene 9 minor version. Since upgrading to Lucene 10 would require EBX® to be compiled with Java 21, the index format is now stabilized, and new releases will not require a full index rebuild in the near future.

EBX® Container Edition

To configure a JNDI data source for EBX®, you'll need to customize your Docker image.

See *Configuring the database in JNDI* for more information.

Closed issues

See [6.2.2 Closed issues](#) section.

Additions to 6.2.1 release notes

The following were included in the EBX® 6.2.1 release, but were not documented in the release notes:

- Complete table revision rework:
 - A new encoding was introduced that improves storage efficiency without the use of compression. This improves speed and puts less demand on resources.

- Updates to the G_SHR database table allow sharing between similar (but different) rows. With the new encoding and sharing detection, a combined size for the G_TRV and G_SHR tables was significantly reduced in a production environment (size reduced from 90GB to less than 800MB).
 - A new in-memory data structure is used to represent table revisions on multiple levels. This supports data sharing to reduce the memory consumption. It is lighter to copy, resulting in faster isolation of procedure and query execution.
 - The checksum computation on index files is faster. This results in faster index opening times by avoiding the reading of files one byte at a time.
 - Updates were made to pooling of the primary key bloom filter on index leaves. Previously, structures used to represent the indexes on the Java heap were not shared, even though they represent the same file on the disk used by several dataspace. A new index format enables sharing this type of file, which consumes the most memory (up to several hundred MB). Although index rebuild is not required when upgrading EBX®, it is required to take advantage of this improvement. In a real-world example, it can save large amounts of memory when the same large file is used in multiple dataspace (in a heap dump, the same 512MB file opened 22 times previously used 11GB of RAM and now only uses 512MB).
 - There was a reduction in the size of the primary key bloom filter on index leaves by increasing its false positive probability. This doesn't have a big impact on performance, but can divide disk and memory consumption by 2 (in most cases).
 - Index writers on indexes that are assumed to be read-only are no longer opened. With fewer objects, memory consumption and pressure on the GC are reduced.
-

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 <i>Dataset trigger</i> 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 <i>Single Sign-On (SSO)</i> 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 <i>SCIM directory</i> 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 <i>SSO SAML 2.0 connectivity</i>, <i>SSO OpenID Connect connectivity</i>, <i>Bearer Authentication</i>, <i>SCIM directory</i>, <i>LDAP connectivity</i> and <i>Directory connectivity</i> 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 <i>Exception handling</i> section was updated with the description for a custom exception mapper. See <i>REST authentication</i> and <i>SOAP authentication</i> 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 <i>Supported databases</i> for more information.

User interface validation	A new option has been added to customize the behavior of APV. <i>See also</i> <i>APV should store invalid input in DataContext</i>
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: • <code>ebx.security.loginPage.appendResumeUrl</code> • <code>ebx.security.exitPage.appendResumeUrl</code> • <code>ebx.security.exitErrorPage.appendResumeUrl</code> See <i>Security configuration</i> for more information.
Security - HTML headers	New options have been added in order to include additional JavaScript files or to customize the CSP. • <code>ebx.security.csp.script-src</code> • <code>ebx.security.csp.style-src</code> • <code>ebx.security.csp.img-src</code> • <code>ebx.html.header.javascript</code> See <i>Security configuration</i> and <i>HTML headers</i> 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, <i>CKEditor</i> was the library used for HTML inputs. For security and licensing reasons, it was replaced by a widget based on the open source tool <i>TipTap</i>. As a consequence, the look and feel has changed. As outlined below, some features were moved or deleted: <i>Look & feel</i> • 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 <i>Ergonomics and layout</i>: 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. See <i>Quick search result highlighting</i> for more information.
Performance on highlights with linked fields	Improved behavior when searching with linked fields to limit highlight calculations on linked tables to only the displayed elements.
Performance during the phase of optimal query plan calculation	Some frequency calculations used by the Calcite optimizer, notably to determine the selectivity of Lucene filters, in order to find the best way to execute the query, have been improved.
New software delivery site	You can now download EBX® software from <i>TIBCO eDelivery</i> . If you have an existing eDelivery account, your existing credentials can be used to login to the new software delivery site.
Closed issues	See 6.2.1 Closed issues section.

Java API

The Java API was improved on the following topics:

- WidgetDisplayContext

New method **WidgetDisplayContext.isInherited()** `WidgetDisplayContext.isInheritedAPI`.

- PerspectiveTheme

New methods **getButtonsHighlightTextColor()** `PerspectiveTheme.getButtonsHighlightTextColorAPI`, **getFormHTMLEditorsMaxHeight()** `PerspectiveTheme.getFormHTMLEditorsMaxHeightAPI`, **getFormHTMLEditorsMaxWidth()** `PerspectiveTheme.getFormHTMLEditorsMaxWidthAPI`, **getFormHTMLEditorsMinHeight()** `PerspectiveTheme.getFormHTMLEditorsMinHeightAPI`, **getFormHTMLEditorsMinWidth()** `PerspectiveTheme.getFormHTMLEditorsMinWidthAPI`, **getPerspectiveNavigationPaneSelectionTextColor()** `PerspectiveTheme.getPerspectiveNavigationPaneSelectionTextColorAPI` and **getWorkflowBadgeTextAndBorderColor()** `PerspectiveTheme.getWorkflowBadgeTextAndBorderColorAPI`.

- ValidationPolicyForTableViewEnum

New enumeration **ValidationPolicyForTableViewEnum** `ValidationPolicyForTableViewEnumAPI`.

- SchemaDefaultTableView

New interface **SchemaDefaultTableView** `SchemaDefaultTableViewAPI`.

- SchemaNode

New method **SchemaNode.getSchemaDefaultTableView()** `SchemaNode.getSchemaDefaultTableViewAPI`.

- 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 seamless 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).

New properties added to replace work item label and description (workItemLabel and workItemDescription) for Work item related tables (Inbox, Work item and in Workflow administration).

New property added to replace Workflow current step (dwfTokenCurrentStep), Workflow description (dwfDescription), Workflow creator (dwfWorkflowCreator), Current user (dwfProcessInstanceCurrentUser), for (Active/Completed workflows and in Workflow administration).

New property added to replace User (allocatedTo), Offered to (offeredTo) for Work item related tables (Inbox, Work item and in Workflow administration).

New property added to replace Comments (dwfComments) for Work item related tables (Inbox, Work item and in Workflow administration). This one is particular, the old property is the one available for display but the one available for filtering is the new property. They share the same label.

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 optimizations

- **Inbox counter optimization:** The workflow inbox counter has been refactored to cache its value more efficiently, significantly reducing database load and improving response time. The new implementation leverages advanced caching mechanisms and optimized queries, resulting in faster loading of the inbox and a more responsive user interface, especially in environments with a large number of workflow tasks.
- **Deadline calculation improvements:** The deadline calculation process has been enhanced to be more efficient. Previously, deadlines were calculated dynamically, requiring recalculation for each deadline at every refresh of the counter. Now, deadlines are calculated at workitem creation and cached appropriately, which reduces system resource usage when many workflows with deadlines are running simultaneously. This optimization eliminates the need for dynamic recalculation, particularly improving performance in environments with complex requirements.
- **Adaptation filters enhancements:** The adaptation filters used in workflow contexts have been optimized to improve performance. Their usage for deadline calculations has been reduced and they are now applied less systematically, as the new batched calculation approach replaces the previous dynamic filtering system.

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 `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.

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 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 *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 `PrefilteredAdaptationFilterAPI` has been added for adding a prefilter to an `AdaptationFilterAPI`.

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:

EBX® Metadata management	It is now possible to build an image including EBX® Metadata components. See <i>Building the image</i> for more information.
Java SE Development Kit	The JDK has been upgraded to LTS version 21. See <i>Overview</i> 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

Building the image

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-assistant.war	This is the WAR for the new AI Assistant. <i>See also New feature: AI Assistant</i>
ebx-ide.war	This is the WAR for the new Integrated Development Environment. <i>See also Overview of the Script IDE</i>
ebx-hub.war	This is the WAR for the new Hub web application. <i>See also Home page</i>

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.

JVM Support	Starting with EBX® version 6.2.3 you can use JRE 25. Use of one of the following Java Runtime Environment is required: JRE 25, 21, or 17 LTS.
Supported databases	Backward compatibility - SAP HANA as a back-end database engine for EBX® is discontinued as of this release; the inter-database migration provides a way to export the full content of a repository to another database. Oracle 21c (Innovative Release) is no longer supported; the required version for Oracle is now 19c or 23c. Microsoft SQL Server 2016 will come out of support in an upcoming release. We advise that you upgrade the database to a recent supported version. See <i>Inter-database migration</i> for more information.
WebLogic AS support	Currently, no version of WebLogic AS supports Jakarta EE 10. 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.

Attention

The upgrade process is not reversible; once a repository is used with a newer version it cannot be downgraded.

The following are high-level steps to upgrade your repository. Please refer to the *Installation Guide* for more environment-specific deployment details. To upgrade your version of EBX®:

1. Shutdown and back up your repository, including the RDBMS and the `ebx.repository.directory` directory. This step is important as it is not possible to rollback a repository.
2. Replace all existing WAR and JAR files on your application server with those from the latest EBX® version. This should also be done for any installed EBX® add-ons.
3. Start your server.

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 dataspace. This simplifies the validation process by avoiding to create snapshots for functional dataspace 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, these 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.getAssociationResultAPI` and `AssociationLink.getTargetTableAPI`

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.

Disable constraint pre-validation

A new method was added to prevent, in the user interface, the constraint from being pre-validated dynamically when the related field is updated at user inputs. This can be useful to improve performance or to avoid unnecessary validation errors during form updates. However, the constraint will still be validated when the form is submitted or when loading the record in the view.

See `ConstraintContext.disableAjaxPrevalidationAPI` and `ConstraintContextOnTable.disableAjaxPrevalidationAPI`

6.2.3 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.

6.2.3 Closed issues

Assistants

CP-28645

The 'temperature' parameter of '0.0' in the LLM configuration generates an error.

Core Engine

CP-26987

A dataspace cannot be deleted after merging it into its parent in the same procedure.

When trying to delete a dataspace that has been merged in the same procedure an exception is raised. The exception indicates that the dataspace to delete has children.

CP-28005

In certain cases, before commit triggers can fail or have inconsistent execution.

issues during before commit trigger execution:

CP-28334

An exception occurs when comparing two PrimaryKey objects from different tables.

CP-28405

An XML parsing exception occurs when using JRE 23 or newer.

CP-28470

Migration order for G_SHR rows during inter-database migration should be enforced.

Inter-database migration fails if a G_SHR row with valued FULL_FK is migrated before the G_SHR rows it refers to. The Inter-database migration process must enforce an order on the G_SHR table to ensure the referenced rows are migrated before the rows that reference them.

CP-28490

The 'archive_id' column of the G_BLK database table is not indexed.

This prohibits foreign key enforcement on row deletion.

CP-28491

A container data model is recompiled after purging a history table that was disabled.

After purging an history table that has been disabled, its former container data model is needlessly recompiled. This can lead to performance issues when performing a batch processing for purging a set of history tables.

CP-28497

Queries on inherited tables do not provide an exhaustive list of involved tables.

In a query on a table in an inherited dataset, the system only uses the table in the root dataset to initialize and isolate the transaction context. As a result, if the inherited table or non-root ancestor are modified in the same procedure, the changes are indexed but are not visible to persistence. This can lead to inconsistencies and can raise issues such as RecordNotFoundExceptions.

CP-28564

When the IO rate limiter is active for Lucene merge threads, it results in a slower merge.

The performance degradation is noticed when creating a child dataspace after a large import.

CP-28591

The processes during a 6.2.0 product upgrade should be split into smaller batches and transactions.

The 6.2.0 product upgrade includes a script that archives oversized blocs ensures that no space is lost in blocs, and computes their last transaction. The script processes all blocs in a single batch and a single transaction. This can reach the limits of some database engines. The process must be split into smaller batches and transactions and allow recovery on partial migrations.

CP-28631

A left over temporary index file must not cause a CorruptIndexException to be thrown.

CP-28647

Preserve the original GrsTableTRV.doUpdateShr() exception in the logs.

If an exception is thrown in the try bloc of GrsTableTRV.doUpdateShr(), it will likely make the release statement also fail and throw an exception that will hide the original. All exceptions should be raised in the logs.

CP-28649

The TableRevision encoding fails on a strictly monotonically decreasing sequence of block references.

The encoded page size must not exceed 32 block references if the ref delta<0.

CP-28714

The PrimaryKeyFilter does not handle foreign keys correctly.

In a table where the primary key is only defined as a reference to a primary key in another table, the internal filter generates an exception.

CP-28728

The index can be corrupted after rollback when pooling is used.

After a rollback, a pooled segment can stay in memory longer than expected. If a new index segment is created with the same name before the garbage collector evicts the outdated segment from the pool, any searcher refresh will use the outdated segment, not the new one.

CP-28789

Boot fails when the database is restored to a previous version and related indexes are not restored.

If an EBX® database is restored from a backup without cleaning the index folder, the transaction discrepancy between table revisions is not detected if the index is ahead of the table revision.

CP-28884	An invalid transaction value for the index can occur in memory cache. In rare conditions, when the pressure on GC his high, the index searcher for the last transaction can be closed but the value of the previous transaction can be used to reopen the index. This can trigger an error like the following: <code>com.orchestranetworks.cce.lucene.index.LxiException: requested transaction 1807732635353088 is in the future of the last transaction 1807732635210753</code>
CP-28900	CteEnumerableJoinRule can generate a join with NULL scores. Depending on the calculation of the estimated rowCount for each side of the join and based on the threshold <code>SMALL_INPUT_SIZE</code> (1000), we will use either a <code>CorrelateJoin</code> or a <code>HashJoin</code>. The <code>CorrelateJoin</code> handles the projection of the filter to the right before the join, which eliminates "residuals", rows potentially with NULL scores. This is not currently done in the <code>HashJoin</code> case.
CP-28932	A ClassCastException can occur when searching for Record Highlights.
CP-28974	A bloc may have been saved in an inconsistent state. A bloc may be stored in an inconsistent state, resulting in an error: <code>"IllegalArgumentException: No free slot found in BlkBloc"</code>. The bloc is marked as having at least one free slot when, in fact, no slot is actually available.
CP-29042	Merge speed should be improved for tables with a composite primary key.
CP-29050	An index exception can occur during or after a purge, which prevents removal of outdated indexes.
CP-29119	Highlight suppliers can retain a lot of memory on subqueries.
CP-29185	Rollback on a table revision without an index can trigger an emergency shutdown.
CP-29223	Connections should be explicitly rolled back after migrating to version 6.2.1.
CP-29231	A nested procedure fails when it unshares the same table revision as its enclosing procedure.
CP-29250	Living transactions are retained for too long and should be removed sooner.
CP-29269	Deleting a temporary file for an invalid directory should not throw an exception.
CP-29436	Index searcher from a nested transaction can be lost in parent procedure. When several nested Procedures having themselves nested procedures are executed under the same root procedure, the index searchers in the refreshed context of the second level-1 sub-procedure are not properly retained in the context of the root procedure. This leads to an <code>AlreadyClosedException</code> if the GC collects them before they are accessed.

Data Modeling

CP-28615	Validation fails when comparing two versions of a model in a module. The error happens if the compared model has defined included schemas in the module and if the model is set in the same folder as the included schemas.
CP-28973	An AlreadyClosedException can occur on the index file.
CP-29129	An xs:pattern facet should not be converted to a facet on-set when defined on fields that are not indexed. An exception is raised when validating a field that has no index and defines an <code>xs:pattern</code> facet. This can occur on inherited fields (those that use <code>osd:inheritance</code> property), or on fields that have a <code>VOID</code> search strategy.

Data Services

CP-28524	REST history category and insert using form-data-compact category aborts and returns a 404 error.
----------	---

Data Validation

CP-28544	The last transaction id can be incorrectly persisted when an error occurs during the validation of a dataset. In some rare cases, technical information used to perform incremental validation is not correctly persisted. As a consequence, an error can occur when the system tries to retrieve the dataset's validation status in the UI. When this happens, no actions are available in the user interface for the related dataset.
CP-28652	An exception can occur when validating a dataspace that has one or more child dataspace. The issue can happen only in a specific scenario when a dataset exists in the parent dataspace, but is missing or in a severe error state in one of the child dataspace.

Data Workflows

- CP-28579 The permission denied log message should be improved and clarified in `WorkflowEngineforceTerminationWithoutCleaning & terminateProcessInstance`.
- CP-28590 The `workflow.lastComment` and `workflow.lastDecision` built-in variables cannot be evaluated for mail with `ScriptTaskSendMail`.
- CP-28927 Workflow profile system variables not working.
System variables `user.login` and `user.fullName` not interpreted.

Distributed Data Delivery

- CP-26896 The broadcast aborts when the replica delivery snapshot is already closed.
- CP-28538 Data services operations in D3 replica nodes may fail during broadcast.

Documentation

- CP-28780 The documentation should declare that setting Tomcat's `unpackWARS` to `false` is not supported.

Hierarchies

- CP-27809 Some actions are not working correctly when using a workflow's hierarchy view.
- CP-28312 A `NullPointerException` occurs on a pruned hierarchy that uses a time dimension.
- CP-28383 A `ClassCastException` can occur when using the 'Find' tool with 'Detect cycle' enabled.
When a hierarchy dependency is based on an incomplete ternary link record and the 'Detect cycle' parameter is set, there is a risk of a `ClassCastException` occurring if the 'Find' tool is used on the hierarchy.
- CP-28602 Selection of recursive member may be long due to cycle detection
On big tables, when a recursive member is selected, it takes long before the number of selected members is displayed.

Infrastructure

- CP-28786 SSO: enhance the JWKS load system and add properties to choose default reference claims and signing algorithm.

Mame

- CP-29160 During the clustering phase, incorrect data is being passed to the EBX® Match and Merge Add-on's matching pipeline.

User Interface

- CP-25689 An error can display in a perspective's validation report after a section was removed from its menu.
- CP-28255 Table header information buttons cannot be accessed by keyboard.
- CP-28344 The screen continuously loads when the browser reaches its iframe maximum depth limit.
- CP-28419 A synchronization issue occurs in the UI.
- CP-28456 When navigating to a table, the sort of the default view may not be applied.
- CP-28473 Light text color is not applied to perspective items.
- CP-28663 It is not possible to scroll on a fixed column in tabular views.
- CP-28740 The `/log/v1` REST service should be removed.

Users & Permissions

- CP-28622 Synchronized directory users are not mapped to their specific roles when calculating permissions.
This can prevent calculating permissions when using LDAP in version 6.2.1 and upgrading to 6.2.2.

CP-28638	An LProjectException is thrown when loading data for a table that has a scripted rule on it.
CP-28666	The LDAP implementation's compatibility with session cookies is not documented.
CP-28673	An error occurs when migrating to LDAP synchronized directory from a version prior to 6.2.1. The synchronized directory should be the first thing installed on a repository since other installations may depend on it.
CP-29094	An LProjectException is thrown when loading data for a table that has a scripted rule applied to it.

6.2.2 Closed issues

Container Edition

CP-28033 **The EBX® GO Add-on should not be considered a legacy add-on by the Container Edition installer.**

Core Engine

CP-26903 **The system should avoid retrieving the cluster center from similarity query results.**
The center of a cluster is not found under certain conditions related to null options. This can result in missing matching results. Not retrieving the cluster centers can also result in a performance improvement.

CP-27719 **A LockObtainFailedException: lock instance already obtained exception can occur.**
In rare occasions when there is a high thread count and low memory, a Lucene LockObtainFailedException can trigger an unwanted index rebuild.

CP-27725 **Running a search on foreign key fields does not display results as expected.**
When searching by keyword, the results are not always sorted logically. Results that correspond to an exact search should be considered the most relevant and displayed first, ahead of the partial results.

CP-27812 **Index files are not always deleted from storage.**
Deleting an index (mostly with purge) will not decrement the shared file reference counter in the parent. This can permanently leave an unused segment file on storage, which can significantly increase the size of the index folder on disk.

CP-27909 **The 'Disable Validation' button is visible for linked fields, but should be hidden.**

CP-27969 **A CustomView error can occur with a predicate filter on single-value foreign keys.**

CP-28003 **The Lucene score value can be missing during a quick search.**
In some cases, the Lucene score is "trimmed" from the query. This results in a NullPointerException during the quick search.

CP-28004 **The IndexReader is closed when querying a table that was truncated.**

CP-28071 **Accessing table history can cause an exception to occur when a UI filter, or custom view associated to a deleted table exists.**

CP-28083 **After publishing a data model that contains validation errors, no error displays.**

CP-28160 **Migration from EBX® version 5 to 6.2.1 throws an ArrayIndexOutOfBoundsException when the last blocs of the table revision are shared.**
Accessing a bloc revision throws the exception when the size of the local bloc container is not be adjusted to the total size of the table revision when the last blocs to add are shared.

CP-28183 **A ProcedureContext doesn't always return the expected record after creation or update.**
It is possible for a trigger to perform the same create or update operation on the same table as the one requested by a ProcedureContext. In this case, the returned record is the one handled by the last trigger instead of the one directly created or modified in the procedure.

CP-28212 **An exception occurs during a search (including highlight) on a list's descendant fields.**

CP-28254 **When an index is being rebuild, the reference pooling thread might block acquisition of the WriteLock.**

CP-28355 **An error occurs when enabling history on table when it has been previously purged from the database mapping section.**

CP-28454 **Computation of the index rebuild chain fails after table truncation.**

Data Modeling

CP-27832 **When publishing a data model, an incorrect warning is produced regarding primary key evolution.**

CP-27935 **The predicate error message is not correct when attempting to publish and a missing node path contains an '_' character.**

CP-28325 **An exception occurs when trying to export or publish a data model in the DMA.**

Data Services

CP-28296 **The Team Up REST API returns HTTP 500 errors when any comment, rating, or tag is added to the record.**

Data Validation

CP-28110 **An error occurs when validating the 'ebx-dataSpaces' dataset in the administration section.**
This error indicates that a procedure is already running on the container datapace.

Data Workflows

CP-1973 [29286] **The 'See more' feature of the work item description inside the workflow inbox does not work.**

This issue happens when users create and launch a work item with a carriage return in the description.

CP-27755 **The hierarchy view in workflows no longer displays.**

CP-28144 **Only active and used workflow publications should be retrieved for better filtering performance.**

Documentation

CP-28056 **Language selection for the documentation should happen automatically based on the browser's locale.**

Hierarchies

CP-27790 **A level filter does not work properly on a hierarchy based on an incomplete ternary join table.**

CP-27886 **A `ClassCastException` can occur when using the Find tool on a hierarchy with a filtered link record.**

When a hierarchy dependency is based on a link record and a filter is applied to that link record, there is a potential risk of a `ClassCastException` occurring if the user attempts to use the `*Find*` tool on the hierarchy.

CP-27904 **When using the Find Tool on a recursive hierarchy, the operation may take an unexpectedly long time if some records in the hierarchy are invalid.**

CP-27966 **Hierarchy errors occur when using an incomplete ternary dependency.**

CP-28220 **A `StackOverflowException` can occur on Hierarchy cycle detection.**

In the case of recursive hierarchies or workflows with deep nesting, cycle detection may trigger the exception.

CP-28230 **The level filter is ignored on a target table that has an incomplete dependency on a ternary link table.**

When a dependency is defined through a ternary link table, and one or more of the foreign keys in the ternary table are missing (i.e., the dependency is incomplete), the `*level` filter on the target table is not applied* as expected.

Infrastructure

CP-27885 **LDAP administration does not work when role display is not configured.**

Scripting

CP-27288 **The input mode value is lost during workflow script task migration, which causes the migration to fail.**

CP-28032 **The value of a complex type under a group remains null when updated by a script.**

User Interface

CP-27711 **When finishing a user task, the next page might not display.**

CP-27718 **After applying a filter, it is not possible to delete records where an association is used.**

CP-27721 **Access is denied at the end of a workflow execution.**

CP-27768	The default locale is not being used for localized fields.
CP-27820	Custom error messages are not displayed on the login page.
CP-27870	Labels in the breadcrumb are double escaped in certain cases.
CP-27883	A copy/paste error occurs while in Grid edit view of an association table.
CP-27936	The custom form editor is not displayed.
CP-28080	The <code>ebx.servlet.useLocalUrl</code> property value is ignored on logout.
CP-28135	When selecting a <code>UITableFilter</code> in the UI, the wrong filter might be checked.
CP-28170	Compilation issues can occur when validating the ebx-views technical dataset. The ebx-views technical dataset contains the customs views defined in the repository. A constraint checks if the data model related to a view is valid and compiles it if it has not been yet used. However, if the data model defines relationships to other datasets, it may be incorrectly marked as invalid. This occurs since these datasets do not exist in the current compilation context, specifically the ebx-views dataspace.
CP-28188	Deletion of a selected record by a custom user service is not handled properly. After executing the deletion service, an empty page displays and a exception is logged.
CP-28236	The redirect URL should be checked when closing a user service.

Users & Permissions

CP-27766	It should not be possible to cancel queries on the <code>DefaultDirectory</code>'s roles table.
----------	--

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 NullPointerException 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 BlkBlk.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 dataspace 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 its 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 <code>ClassCastException</code> 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 dataspace. 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 <code>PrimaryKeyException</code> occurs when deleting a hierarchy member. In a ragged hierarchy, a <code>PrimaryKeyException</code> 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 <code>ebxFile</code> appender is specified. The INFO threshold is used instead.
CP-27571	The <code>DirectoryHandler</code> 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 <code>NullPointerException</code> .
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 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.
----------	--