

Administrator Guide

- [Overview](#)
- [Quick Start](#)
 - [Configuration process map](#)
 - [Task 1: Updating the model](#)
 - [Task 2: Registering tables](#)
 - [Task 3: Setting up matching](#)
 - [Task 4: Configuring a merge policy](#)
 - [Task 5: Activating matching](#)
- [Matching with business objects](#)
- [Creating matching views](#)
- [Excluding records](#)
- [Modifying merged records](#)
- [Validating matching configurations](#)
- [Setting trusted sources](#)
- [Matching with relationships](#)
- [Understanding matching operation processing](#)
- [Perspectives and workflows](#)
- [Displaying matching metadata](#)
- [Matching algorithms](#)
- [Performance recommendations](#)
- [Replicating matching metadata](#)
- [Scheduling matching metadata purges](#)

Overview

In an administrative role with the TIBCO EBX® Match and Merge Add-on, you can perform actions that:

- Specify the data that can participate in the add-on's matching and merging operations.
- Create a matching policy that determines *when* and *how* matching processes execute. How you enable add-on features for business users depends on use case requirements. To prevent duplication of existing data, you might want to trigger matching when a business user creates or updates a record. Alternatively, you might have existing data that you want to de-duplicate.
- Create a set of rules that determines how the add-on automatically merges data to create a *golden record*. A golden record contains the most accurate and up-to-date data.

The *Administrator Guide* aims to explain how to configure the add-on. The higher-level administrative tasks required to get up and running are described in the *Quick Start* section. Its chapters and sections are ordered in the recommended sequence of completion. The other chapters and sections in this guide drill into more detailed subjects. After you become more familiar with the add-on, feel free to use this guide as a reference for configuration-related tasks.

What's next? Familiarize yourself with the add-on configuration process using the information in [Configuration process map](#). Or, feel free to jump into [Task 1: Updating the model](#) to begin making your data available to use with the add-on.

Quick Start

- [Configuration process map](#)
- [Task 1: Updating the model](#)
- [Task 2: Registering tables](#)
- [Task 3: Setting up matching](#)
- [Task 4: Configuring a merge policy](#)
- [Task 5: Activating matching](#)

Configuration process map



(/content/dam/onebx/us/en/ebx-addon/mame/6-2-3/admin_guide/resources/images/config process map 1.png)

The following table gives you a high-level outline of the above configuration tasks:


1) Make data model updates	Before using add-on features, you must update your data model .
----------------------------	---

2) Register tables	You make data available to the add-on by registering tables . You can also specify whether to pause add-on functionality on registered tables.
3) Set matching options	You configure matching behavior for a registered table by adjusting options that set <i>how</i> and <i>when</i> matching executes. Matching operations locate potentially duplicate information.
4) Configure a merge policy	You create a merge policy to tell the add-on how it should handle matching results. For example, when duplicate information is found, these settings are used to specify which values should be merged, or survived, into the golden record.
5) Activating matching	When you activate matching configurations the add-on automatically validates configuration settings and informs you of any errors or warnings.


Task 1: Updating the model

Before creating a match configuration, you must update and publish the EBX® data model using the Data Model Assistant (DMA). The steps below describe how to include add-on functionality in your data model:


1. From the main menu bar, select the **Data models** panel and select the data model to update.



2. Expand **Configuration** and select **Add-ons**.



3. Select '+' to add a new entry and choose **Match and Merge**.




4. After saving, publish the data model.

What's next? Follow the instructions in [Task 2: Registering tables](#) to enable add-on functionality in a table.


Task 2: Registering tables

By registering and activating a table, you make its data available for match and merge functionality. To register a table and activate matching:

1. From the main menu bar, select the **Administration** panel and open *Data Quality & Analytics > TIBCO EBX® Match and Merge Add-on > Table activation and settings*.




2. As shown above, select the '+' icon to begin a new configuration.
3. Specify the table to register using the **Data model** and **Table** properties.



4. Select **Save** and do not close at this time.

Several new tabs display, including: 1) **Matching policies** and 2) **Merge policies**. See the steps in the following sections to configure these options: [Task 3: Setting up matching](#) and [Task 4: Configuring a merge policy](#)

If you expand the **Advanced settings** group, you can specify whether users can modify records that were included in a merge operation. See [Modifying merged records](#) for more information on this option.



What's next? Follow the instructions in [Setting up matching](#) to configure a matching policy.

Task 3: Setting up matching

- [Task 3: Overview](#)
- [Task 3a: Creating a matching policy](#)
- [Task 3b: Creating a decision tree](#)
- [Task 3c: Configuring field specific settings](#)

Task 3: Overview

A matching operation compares data from two or more records to find potential duplicates. You create a *Matching policy* to tell the add-on:

- *when* to execute a matching operation. You can require users to manually execute a matching operation. Alternatively, a matching operation can run automatically when users create or update records.
- *how* to execute a matching operation. You use a *decision tree* to provide the add-on with criteria that determines whether records are duplicates. You can also apply a merge policy that takes action based on the results of the record evaluation. Please note that creation of a merge policy is covered in [Task 4: Configuring a merge policy](#).


What's next? Follow the instructions in [Task 3a](#) to configure a matching policy.

Task 3a: Creating a matching policy


You create a matching policy under and apply it to a table registered with the add-on. When creating the matching policy you tell the add-on when to execute. Users can run matching manually from a table menu, or you can specify that matching executes inline during record creation or update. When matching executes inline, a matching policy automatically executes after creating or modifying the record. If the matching operation finds potentially duplicate (suspect) records, users are notified and can review and make changes to the record.

To create a matching policy (skip to step 2 if picking up from [Task 1](#)):

1. From the main menu bar, select the **Administration** panel and open *Data Quality & Analytics > TIBCO EBX® Match and Merge Add-on* and open the table settings you want to edit.
2. Select the **Matching policies** tab.



3. Select the '+' icon to create a new matching policy and fill in the required fields on the page that displays. The image below highlights required options:



Matching policy code	The unique code used to identify this matching policy.
Active	Sets whether this matching policy is active.
Execute matching group	The Execute matching group determines how matching executes during creation or update of records. If set to:

- **Manually** the add-on does nothing and users must initiate matching operations. They can do this from a table's **Actions** menu.
- **After submission** the add-on automatically executes matching after creating or updating the record. You can apply this setting when you want to perform bulk creation of records using the EBX® API. There are limitations to this functionality. Please check the *Known issues* section in the release notes for additional details.
- **Before creation** users can search for potential duplicates before creating the record.
- **Inline** the add-on automatically executes matching during the creation/update process. The user must review any potential duplicate records before continuing. When specifying that matching execute inline, the add-on also displays the following menus to select a workflow for record creation or update: **Workflow on creation** and **Workflow on update**. The selected workflows are launched when a user takes the corresponding action. Note that depending on the assignments in the workflow, the user who created the record that triggered inline matching might not see the screen requiring that the records be reviewed before continuing.

Note:

	If you set matching execution to After submission, Before creation or Inline, you must set this matching policy as the table's Default matching policy. This option is available on the configuration's Main tab. Additionally, when a user creates or updates a record, the add-on will only execute automatically if a field configured for matching is modified.
Execute purge	Determines whether the add-on automatically removes old matching results. If set to Automatically the add-on executes this policy each time matching executes. As this process can be time consuming, you can set it to Manually and allow users to execute as needed. They can access this functionality from a table's Actions menu by selecting: <i>Match and Merge > Purge old results</i>.
Merge policy	Tells the add-on which merge policy to link to this matching policy. After creating a merge policy, you must specify it here in order to activate the automatic merge feature. See Task 4: Configuring a merge policy for instructions on creating a merge policy.
Show match type	Sets display of the Match type column when users manage groups of matched records. This column indicates whether each record was included in the group through a direct match, or via transitive logic. The group's survivor record is identified with an N/A label. When this option is enabled, a built-in filter in the Manage group screen also

allows users to display only records that are transitive matches.

Transitive matching refers to the inclusion of records based on indirect relationships. For example, if Record A matches Record B, and Record B matches Record C, then Record A and C are grouped together due to their transitive connection through B.

1. After saving, the following tabs display:

- a. **Decision tree:** Contains settings that determine how the add-on compares data values.
- b. **Matching fields:** Contains configuration settings for each field used in a matching policy.


What's Next? Use the instructions in [Task 3b: Creating a decision tree](#) to create a matching decision tree to specify *how* this matching policy compares data values.

Task 3b: Creating a decision tree


You create a matching decision tree to tell the add-on which data values to compare and how to compare them. Based on the comparison result you can specify that the add-on set a record's relationship status, or perform additional tests to ensure accurate results. A relationship status identifies records as positive or negative matches, or as suspected of being a match. The add-on creates groups of records identified as positive matches.

To create a decision tree:

1. Select the **Decision tree** tab.



2. As shown below, 1) drag a **Data comparison** test from the **Toolbox** and 2) drag the bottom connection point from the **Start** node and link it with the test.



3. Double-click the test to edit its functions:

- a. The **Name** property sets a name for this test that displays on the tree.
- b. **Evaluation method**: the add-on uses the selected method to compare the fields you add.

When you choose an option from the menu, you can select the preview icon to view a detailed description. The following comparison methods are available:

- **All fields match**: For this data comparison to evaluate as true, each included field must meet or exceed its minimum score setting. See the definition of **Minimum score** in step 3 for more information.
- **At least N fields match**: This option allows you to specify that a certain number of included fields must meet or exceed their minimum score setting for this data comparison to evaluate as true. When you select this option, the **Minimum number of matches** property displays and allows you specify how many fields must match.
- **Average**: An average score is computed from the score of each included field. If this score meets or exceeds the value set in **Minimum score**, this data comparison evaluates to true. This function also allows you to assign a weight to each field. When using multiple fields as part of one data comparison, the **Weight** property is used to compute a weighted average and assign a weight to each field. Acceptable values are between 0 and 100. The higher the value, the more weight the field carries.
- **Predicate**: This option is deprecated. Instead, in the **Compare using** field, select **Operator** and the **Equals** option.

c. **Fields** (One or more fields used in the comparison):

- The **Use relationships** menu allows you to specify whether this field is: a regular field, a relationship through a link table, or a foreign key relationship. Before a related field is available to select, you must configure it in the **Matching fields** tab. See [Matching with relationships](#) for more information.
- The **Compare using** option allows you to choose whether to use an algorithm or an operator when comparing the selected field's values. When you select:
 - **Algorithm**: Use the **Algorithm** menu to select the algorithm used to compute a score for this field. The list of available algorithms is filtered depending on the search strategy specified for this field in the data model. In some cases, an algorithm might be available but not recommended for use with a particular search strategy. These algorithms display in a lighter color font at the bottom of the list. If you select an algorithm that is not compatible with the defined search strategy, a

warning icon displays. When configuring the algorithm for an added field using the **Matching fields** tab (see [Task 3c: Configuring field specific settings](#)), a link displays to allow you to change the search strategy defined for this field. After changing the strategy, the list of available algorithms automatically updates.

When you select the:

- **Exact** algorithm, you can specify whether to use case sensitivity on fields with String and Text data types.

Note:

If you select a foreign key field, you should only select the **Exact** algorithm.

- **Hybrid fuzzy** algorithm, you can adjust parameters related to several of the criteria used in the algorithm. The higher you set the value for a particular criterion, the more it will count when calculating the final similarity score. See [Matching algorithms](#) for a more in-depth description of this algorithm.

The **Advanced settings** group allows you to configure how a field's null values are managed in this comparison node. Since the same field can be used in other comparison nodes, you can apply different behavior to the field in other parts of the decision tree. For example, one node can require an exact match, where both records must contain a value—use **No match** for both null value settings. Another node can use a "not mismatched" strategy, which considers records matching unless their values are explicitly different. In this case, set one or both settings to **Match**.

Please note that these null value settings pertain to decision tree evaluation during processing only. Whereas, the null value management options on **Matching fields** are applied in pre-processing during clustering.

- **Operator:** Use the **Operator** menu to select the operator used to compare values in the selected field. If you select **Equals**, a field displays that allows you to enter a value that must be present in the field to consider it a match.

The operators available to select depend on the data type of the chosen field. If the data type is:

- `date`, `time`, `dateTime`, `decimal`, `text`, `string`, `integer`, `boolean`, or `email`; the available options are: **Equals**, **Either null**, **Both null**, and **Both not null**.
- any other data type, the available options are: **Either null**, **Both null**, and **Both not null**.

Attention:

If you use any operator that checks for null values, set the weight for the matching field to 0 in the **Matching fields** tab.


- **Minimum score:** Sets the minimum percentage score required to consider data values from this field as matching. Acceptable values are between 1 and 100. The higher you set this value, the more similar the data values must be to match.

Note:


You cannot select computed, multi-value, or linked fields in the **Fields** menu.

4. Specify the relationship status if the test evaluates to *true* by dragging one of the following from the **Toolbox** and connecting it with the test's right output point:

- **Match:** the records will be considered a match.
- **Suspect:** the records are suspected of matching, but require human review.
- **No Match:** the records do not match.




- Specify the relationship status if the test evaluates to *false* by connecting the desired status to the test's left output point. Note that if you want the result to trigger a new test, add another **Data comparison** node and connect it to desired output.



- Use the process shown in the previous steps to build out your matching tree. Once finished, select **Save and close**.

The following example shows using three comparisons. If the first comparison evaluates to true, the records are considered suspects from that point on. If the second and third comparisons also evaluate to true, the records are marked as matches.



7. If required, set the configuration's default matching policy:

- From the **Matching policies** tab:
 - a. Select the policy you want to set as the default.
 - b. From the **Actions** menu, select **Set as default**.
 - c. Save your changes.
- From the **Main** tab:
 - a. In the **Default matching policy** menu select the policy you just created.
 - b. Save your changes.

What's next? Optionally, [configure field specific settings](#). Or, see [Task 4: Configuring a merge policy](#) for instructions on setting up the add-on to merge data.

Task 3c: Configuring field specific settings

Use the matching policy's **Matching fields** tab to access additional functionality related to fields, including:

- Use of relationships in matching.
- Exclusion of records from a matching operation. When configured correctly, record exclusion can improve overall performance. See [Excluding records](#) for more information.
- Pre-processing settings for field:

- **Search strategies:** The search strategy is used to group similar records during the pre-clustering process. The list only includes search strategies defined in the data model containing the selected field. See the EBX® Reference Manual for more information about search strategies.
- **Weight:** The weight is assigned to this field during pre-processing by EBX®. When a matching operation includes multiple fields, the weight sets the relative importance of each field's score in determining whether records are similar. Note that this weight does not affect the weighted average in the decision tree comparison nodes. You might set a higher weight on a field that could determine the output earlier in the decision tree. If you set the weight to 0, this field is excluded from the matching pre-processing phase.

Note:

To improve matching performance, set the weight to 0 for fields that have a limited number of options. For example, fields that are boolean or enumerations. However, at least one or more fields should have a weight above 0 defined. If all fields are set to 0, matching results are not reliable.

By default, during the pre-processing phase of matching, field weights are automatically balanced so that each field equally affects the final score. This means that fields that contain multiple values have the same impact on the final matching score as fields with single values. To disable this behavior, set the `ebx.clustering.balance` property to `false` in the `ebx.properties` file. See [Understanding matching operation processing](#) for more information.

- Null value management determines the behavior when comparing two records and one or both of them has a Null value in this field. When enabling matching of null values, adjust the weight for the field so that it is lower as compared to other fields. Otherwise, it might not get included during pre-processing with other similar records.

What's next? See [Task 4: Configuring a merge policy](#) for instructions on setting up the add-on to merge data.

Task 4: Configuring a merge policy

- [Overview](#)
- [Creating Merge policy settings](#)

- Configuring survivorship fields
- Managing relationships
- Enabling automatic merge

Overview


When a matching operation positively identifies duplicate records, or a user manually initiates a merge, the merge policy determines:

- Which record should *survive* the merge process. In other words, the record into which values will be merged to create a golden record.
- The function used to decide which data values will be included in the golden record. One function can be applied to the entire table, or you can set a separate function for each field.
- Whether a new golden record is automatically created. If not, values are merged into an existing record.
- How FK relationships are handled as part of a merge. For example aligning foreign keys to the new golden record by updating the foreign keys in related tables to target the new golden record.


Creating Merge policy settings

To create a Merge policy (Begin with Step 2 if you completed the previous quick start task):

1. From the main menu bar, select the **Administration** panel and open *Data Quality & Analytics > TIBCO EBX® Match and Merge Add-on* and open the table settings to which you want to add a Merge policy.
2. Select the **Merge policies** tab.



3. Select '+' icon to create a new merge policy and fill in the required fields on the page that displays. The image below highlights these options:



1) Merge policy code

A name used to describe this merge policy.

2) Auto create new golden

Allows you to specify that a new record will be automatically created as part of the merge process. The **Mode** drop-down menu

	allows you to activate and specify when a new golden is created. When this option is enabled, you can use the Customize source value for new golden option (not shown) to specify the source to use. The field populates with trusted sources configured in the following location: <i>TIBCO EBX® Match and Merge Add-on > Reference data > Trusted sources</i>. Note: If the selected source field is a foreign key to another table, its data type must match that of the primary key in the lookup table.
3) Survivorship strategy	Determines how values are survived into new records. Record selection function: Specifies the record selection function to use. A record selection function determines which record survives the merge process. For example, with the Most trusted source option, the record stemming from source with the highest level of trust will survive the merge process. Data values will be merged into the surviving record to create a golden record. Field merge function: You can use this property to specify a single function that determines which data values are merged into the golden record. If you would like to create settings for different fields, use the options in the Survivor field tab to do so.
4) Used for manual merge	Sets whether this merge policy's settings are applied when a user manually initiates a merge operation. If multiple merge policies

are defined for the same table, you can only enable this option on one policy. When enabled:

- Values for the manual merge are pre-selected based on this merge policy's settings. Users can still update selections from the merge screen.
- The **Apply permission on merge view** field displays. This option allows you to set whether TIBCO EBX® permissions apply in the **Merge view**. This allows you to address the use case where a data steward needs to access all fields in order to perform a successful merge.

Configuring survivorship fields

Merge policies allow you to define criteria for entire tables, but also allow you to set parameters for individual fields. To set merge behavior for a specific field:

1. Select a merge policy's **Survivor field** tab and use the '+' icon to set up a new field.
2. The following list highlights some of the options:
 - a. **Field**: Choose the field on which these settings apply.
 - b. **Field merge function**: Select the function used to choose which value should survive the merge process. If you choose **Constant**, you can enter a value that the add-on will populate this field with (in all golden records). Note that this option is only available for fields of String and Text data types and *not* for multi-valued fields.
 - c. **Condition for field value survivorship**: Optionally, enter a predicate expression to evaluate the value to be survived. Both the **Field merge function** and this expression must evaluate to true to be survived. The XPath predicate expression used here must comply with EBX® requirements. For more information, see the *Supported XPath syntax* section *EBX® Reference Guide*.
 - d. **Execute only if empty**: Set whether the merge only happens when the value in the survived record is empty.

Managing relationships

The add-on allows you to set behavior for cases when a merge operation involves relationships. For example, you might want to update a foreign key that points to the value in a merged record to point to the value in the golden record. To manage how the add-on handles relationships in these types of scenarios:

Attention:

The following steps apply to related tables contained in the same dataset. If the related table is in a different dataset or dataspace, the add-on will automatically align the foreign keys.

1. Select a merge policy's **Related tables** tab.

The add-on automatically adds options here for related tables that come from the same data model and hold a foreign key to the table where matching executes. Note that if the data model was updated, you can refresh the list of related tables. To do this, open the **Actions** menu and select **Refresh**.

2. Double-click to open the relationship you want to update.
3. Use the **Align foreign keys** drop-down menu to set how FKs should be managed.


You can choose from the following options:

- **None:** This relationship is ignored.
- **Manually:** The merge process will allow users to manually select whether the related value is updated.
- **Automatically:** The FK value is updated automatically based on the survived record.

Enabling automatic merge

If you want the add-on to automatically merge records that it considers positive matches, you must associate a merge policy with a matching policy. Follow the instructions below to edit a matching policy and associate a merge policy:


1. Select a registered table's **Matching policies** tab and open the desired matching policy.



- On the **Main** tab, use the **Merge policy** menu to select the desired merge policy.

Note:

When you select a merge policy here, merging happens automatically and happens according to the settings in the selected policy. When a merge is triggered, it bypasses the option to merge manually.



- Save and close.

What's next? To enable automatic matching features for your data model, see the instructions in [Task 5: Activating matching](#).

Task 5: Activating matching

The final step to setting up a configuration is to active matching on the table configuration. You only need to perform this step once per table configuration. As a best practice you can wait until creating several table configurations and enable them all at once.


Note:

If matching is enabled for multiple tables from the same data model and you need to re-activate one table, you must select and activate all of the configurations from the same model.

To enable matching:

1. Navigate to **Table activation and settings** and select each configuration you want to activate. As a shortcut to activating all configurations, do not select any records.
2. Select **Activate matching** at the top of the screen.

Part of this operation includes validating the matching configuration. You will see a validation page with any errors or warnings that you might want to address prior to using the matching configuration.



After activating matching you can use the add-on's **Evaluate matching** service. This service allows you to run a match on selected record and review the results to ensure they are consistent with your expectations. See [Evaluating matching](#) for more information.

Matching with business objects

- [Overview](#)
- [Business object example](#)
- [Configuring matching on business objects](#)
 - [Automatic golden creation for business objects](#)

Overview

In some business use cases, data deduplication is performed on the data that defines a business concept or object. Often, multiple tables are used to create or model a business object and define its attributes. In these cases, matching needs to be performed in a more holistic fashion to determine whether two objects represent the same entity.

Merging and golden record creation should also take into account the object as a whole. To meet these requirements, the TIBCO EBX® Match and Merge Add-on allows you to configure and run matching and merging operations on EBX® Business Objects.

Note:

The add-on supports matching up to a depth of 3 levels (parent, child, grandchild) on business objects. However, to avoid performance issues, it is recommended to configure a maximum of 2 levels for matching.

In essence, a business object is comprised of one main table, and all of its related tables. See the TIBCO EBX® documentation for more details on business objects.

Note:

The current relationship depth that is supported for merging is the grandchild level. In other words, two relationships away from the main business object table.

Business object example

A common use case for business objects is that of representing a customer. For instance, a main customer table that represents an individual or business. It is possible that many tables are related to this main table and together they comprise the customer entity, for example address, contact, and billing information. From a data deduplication standpoint, two business objects might be considered as representing the same entity when a match is found in the:

- Main business object table.
- Business object's related tables.
- Main and related tables of the business object.

If a match is found, the add-on maintains any relationships. For example:

- When a new golden record is created in the main table, new golden records are also created in the related tables and linked to the new main golden.
- If a record in the main table is determined to be golden, and no conflicting policies are applied to the related tables, all linked records in related tables are also considered golden.
- Since a merge operation will only result in the creation of new golden records, merged records from the main and related tables remain unchanged.

Configuring matching on business objects


Before setting up a matching configuration for a business object, the business object must be created in the DMA and published. Once these steps are complete, you configure matching for a business object much in the same way as for a normal table.

Note:

A table cannot be activated for a business object that is already used in a single table matching configuration.

The main differences in business object matching configuration are:

- You start from the **Business object activation and settings** in the **Administration** pane and register a business object with the add-on instead of a table. Note that the **Advanced settings** tab is located at this level.



- You have the option of creating one configuration for the main table, or configuring each table in the business object.
- When you want to match and automatically merge records in the business object's related tables, you must configure matching and merging on the main table and related tables. Specifically, each related table must have a default matching policy that specifies a merge policy.
- The option to automatically trigger matching operations is available for business objects. Only the **After submission** option is available to use with business objects, either **On creation** or **On update**.

Configure this option on the matching policy for the business object's main table (the option is not available on related tables). When users create records in tables related to the business object's main table, the **On update** property setting is used.

- When configuring decision trees, you can select fields to compare from all tables that make up the business object.
- Matching operations are only available to run from the main table in the business object.
- The following limitations are in place for business object matching:
 - When configuring the decision tree for related tables, there is no **Suspect** node.
 - Only automatic foreign key alignment is available.

Automatic golden creation for business objects

When configuring automatic creation of golden records in a merge policy for a business object, you can choose from the following options:

- **Duplicates and Singletons** (available for business object parent and child entities): The add-on automatically creates a new golden record when it identifies a positive match between records, or when a record doesn't have any match. This property is only applied in the case of record creation, modification, merge or matching operations.
- **Only duplicates** (available for a business object parent entity only): A new golden record is automatically created only when the add-on identifies a positive match between records. With this option selected, child entities must use the **On duplicate parent** option described below.
- **Any duplicates** (available for child entities only): In a Business Object if the parent has positive matches, a new golden is created for the parent and its children. If matches are found for children of the same parent, but not the parent, a new golden is created for the child only.

Creating matching views

- [Overview](#)
- [Creating a tab that shows merged records](#)
- [Creating a hierarchical view of golden records](#)

Overview

You can create matching focused views that include:

- Adding a tab to a golden record's details view that shows all source (merged) records of the golden.
- Creating a hierarchical view where parent and child nodes are golden and merged records, respectively.

Creating a tab that shows merged records

Use the following steps to add an association to your data model using an XPath predicate:

1. Select **Data models** from the main EBX® toolbar.
2. Select the desired data model.
3. Using the ellipsis menu next to the desired table, select **Create child**.
4. Fill in the **Name** field and for the **Kind of element** option, select **Association**.
5. In the **Definition** box, select the following:
 - **By XPath predicate**

- **Target table:** Select the current table
- **XPath expression:** Enter the following `osd:is-not-null(./Linked_Field_GoldenId)` and `./Linked_Field_GoldenId = ${../Linked_Field_FunctionalId}`

6. Select **Create**.

7. Publish your data model.

After successful publication, the details view for golden records is updated with a **Main** tab and an association tab that displays merged records.

Creating a hierarchical view of golden records

You can create a hierarchical view with golden records as parents and corresponding merged records as children.

To create the hierarchy view:

1. Open the table you want to create the view of.
2. From the **View** menu, select **Create a new view**.
3. Define desired options, but for the **View mode**, select **Hierarchical view**.
4. Select **Next**.
5. In the **Dimension** box, select the **Target record** option.
6. Select **Create** to create the view, or **Create and apply** to create and display the view.

This view is now accessible from the **View** menu.

Excluding records

- [Overview](#)
- [Excluding records in a table](#)
- [Business object record exclusion](#)
 - [Prevent business object matching participation](#)
 - [Exclude related records in matching](#)
 - [Excluding records in a business object](#)

Overview

One way that you can boost matching performance is by excluding records from matching operations. When the total number of records the add-on must process is reduced, it can cut down on the total processing time and help improve the overall quality of results by preventing false positive matches. When you configure options to exclude records:

- You tell the add-on to look at the values in a specific field.
- You specify criterion that determines whether (based on the field value) a record is excluded. For example, you could specify that null values in an Expiration Date field exclude records from matching. The reason being, only products that still have an expiration date are not yet expired and should be matched.
- The configuration applies to all available types of matching.
- Any record that is excluded from matching is set to `<Unset>`.

The following sections provide instructions for:

- [Single table configurations](#)
- [Business object configurations](#)

Excluding records in a table

To exclude records from matching enabled tables:

1. Navigate to *Administration > TIBCO EBX® Match and Merge Add-on > Table activation and settings* and open the configuration where you want to exclude records.
2. Select the **Matching policies** tab and open the desired matching policy.
3. Select the **Matching fields** tab and create a new, or open an existing field.
4. Use the **Exclude records where** property to add at least one criteria used to exclude records.

You can select:

- **equals** and **is different** to enter a specific value that must be the same or different to exclude a record, respectively.
- **is null** or **is not null** to specify that if the field is missing a value or has a value the record will be excluded, respectively.

5. Save your changes before exiting.

Business object record exclusion

You can configure record exclusion settings for business objects to:

- Prevent the entire business object from participating in matching operations.
- Stop the business object's related records from impacting the matching outcome.

For the steps to configure business object exclusion, see [Excluding records in a business object](#)

Attention:

When multiple exclusion conditions are configured for the same matching field, the OR operator will be applied.

Prevent business object matching participation

When a record from the business object's main table is excluded, the entire business object is prevented from participating in matching. As a result, the records are not merged; the state for all records that comprise the business object is <unset>.

In the main table's matching policy, the **Exclude records where** setting applies to the main table records whether used on a matching field in the main, or a related tables. This means that you can specify that a record in the main table is excluded based on its values, or on values in related records.

The following image provides an example of excluding business objects from matching operations:

Note:

Results vary depending on the configurations of your match and merge policies.

Main table:		Related table:		
PK	Value	PK	FK	Value
M1	Trang	R1	M1	Hanoi
M2	Trang	R2	M2	Hanoi
		R3	M1	Sai Gon
		R4	M2	Paris
Matching results = Both business objects excluded				
If 'Exclude records where' is configured in the main table's matching policy on a matching field located in the: 1) Main table and set to exclude 'Trang', or 2) Related table and set to exclude 'Hanoi'		Main table: All records <unset>		
		Related table: All records <unset>		

Exclude related records in matching

You can ensure that related records that contain specific values have no impact on the outcome of matching. However, if a match is still found between business objects, the excluded related records are merged and maintain a relationship with any new golden records.

In the main table's matching policy, use the **Exclude related records where** property on a matching field from a related table to define the criteria that exclude records in related tables. Note, this option is only available when both of the following conditions are met:

- The matching policy is configured on the business object's main table.
- The matching field is from a related table in the business object.

The following image provides an example of excluding related records:

Note:

Results vary depending on the configurations of your match and merge policies.

Main table:		Related table:		
PK	Value	PK	FK	Value
M1	Trang	R1	M1	Hanoi
M2	Trang	R2	M2	Hanoi
		R3	M1	Sai Gon
		R4	M2	Paris
Configuration:		Matching results for M1 and M2		
When Exclude related records where is configured in the main table's matching policy on a matching field located in the related table. Value to exclude = Hanoi		If the decision tree is configured so that a positive match is achieved only when the records in the main table AND related table match. M1 Trang <—> M2 Trang: Positive Match R1 & R2 Hanoi: Excluded R3 Sai Gon <—> R4 Paris: No Match Result: No match between business objects.		

Excluding records in a business object

To exclude records from matching in a business object:

1. Navigate to *Administration > TIBCO EBX® Match and Merge Add-on > Business object activation and settings* and open the configuration where you want to exclude records.
2. Select the **Matching policies** tab and open the desired matching policy.
3. Select the **Matching fields** tab and create a new, or open an existing field.
4. Use the **Exclude records where** and **Exclude related records where** properties to determine what values will prevent a record from participating in matching.

See the sections above for examples of record exclusion behavior.

5. Set the record exclusion criteria:

You can select:

- **equals** and **is different** to enter a specific value that must be the same or different to exclude a record, respectively.
- **is null** or **is not null** to specify that if the field is missing a value or has a value the record will be excluded, respectively.


6. Save your changes before exiting.

Modifying merged records

Some use cases require a record's life-cycle to continue after a merge operation where it did not become the golden record. The add-on allows you to enable merged record modification for these types of cases.

To enable merged record modification:

1. Navigate to *Administration > TIBCO EBX® Match and Merge Add-on > Table activation and settings* and open the configuration where you want to enable merged record modification.
2. Expand the **Advanced settings** group on the **Main** tab.
3. Set **Modify merged records** to **Yes**. and choose one of the following options:



4. Determine the add-on's behavior after a merged record is modified by selecting one of the following options:
 - **Keep in group**: After modifying a merged record, it might contain the most accurate values that should be incorporated into the golden record. When you select this option and a merge policy is associated with the default matching policy, the add-on initiates a merge using the associated policy to update the golden record. When no default merge policy is available, the value is updated for the modified record only. The records will stay in their group regardless of whether the modified fields are configured to take part in matching operations.

We recommend this option when modifications do not impact fields configured for matching. This way modifications update the golden record, but do not create conflicts with previous matching operation results.

- **Change state to <unset>**: The add-on changes the modified merged record to unset and removes it from its group. This might result in requiring stewardship tasks, such as running a match and updating any impacted groups. Records are moved out of their groups regardless of whether the modified fields are configured to take part in matching operations.

We recommend this option when modifications impact fields configured for matching. You should re-execute a matching operation to ensure the record is still similar enough with others in its group after modification to consider it a match.

- **Both**: If the record's modified fields are configured to take part in matching operations, the corresponding record will be changed to the <unset> state and moved out of the group. If the record's modified fields are *not* configured to take part in matching operations, the group is remerged based on the default merge policy. When no default merge policy is available, the records are changed to the <unset> state and moved out of the group.

5. Save before exiting.

Users are now able to save changes to merged records.

Validating matching configurations

Incomplete or invalid matching configurations can lead to tainted results and poor performance. To help prevent these issues, the add-on can check configurations and provide a report that highlights any errors or warnings.

Note:

Just because a validation report turns up issues, it will not prevent the add-on from executing an enabled and active matching policy. If the report turns up serious issues, you might want to pause the matching configuration until you resolve them.

To validate a matching configuration:

1. Navigate to **Table activation and settings**.
2. Select one or more configuration entries.

3. Run the validation:

- From the **Actions** menu, select **Validate configuration**.
- Select **Activate matching** at the top of the screen. Note that this option will enable any configuration instances that are not already enabled. However, running this on a configuration that is enabled has no effect other than validating the configuration.

Setting trusted sources

- [Overview](#)
- [Adding and prioritizing sources](#)
- [Data steward decisions as sources](#)

Overview

When you have records that stem from multiple systems, you might trust values from some sources over others. For example, take three source systems: HR, Accounting, and IT. If you had a question about an employee's job title or role, which system would provide the most accurate information? You might even determine that a different source is more reliable for each field. In these situations you can use the add-on's trusted sources feature to specify which source provides the most reliable information for a given field.

The add-on ships with a built-in source called **Data Steward decision**. This source enables users to override the order of trusted sources on a per field basis. See [Data steward decisions as sources](#).


Adding and prioritizing sources

You can create multiple sources within the add-on and define a prioritized list for tables and fields that are included in a merge policy. For this process to work, you need a field in your record that specifies its source. When you use the steps below to add sources, ensure that the sources you add correspond to the allowable values for the field. For example, if a System field had the allowable values of HR, ACCT, and IT, you would create the following sources: HR, ACCT, IT.

To add sources and set their priority:

1. Create the sources:
 - a. Open the **Administration** panel and navigate to *TIBCO EBX® Match and Merge Add-on > Reference data > Trusted sources > Sources*.


b. Create a new record.



c. Fill in the required name and optional description.

Attention:


Source names must be identical to the corresponding values that identify sources in records. For example, if a System field identified its originating source using HR, ACCT, or IT, you must provide these same values for each source name.



d. Save and close the record.

2. Set and prioritize sources for fields or tables:

- a. Select either **For tables** or **For fields** depending on whether you want to set sources at the table or field level.
- b. Create a new record. This example uses fields.
- c. Select the desired field or table and create the list of its trusted sources as described below:




1) Matching table and Field menus	Specify the table or field that this list of trusted sources applies to. Note that only fields of the String and Text data types are available to select from the menu.
2) Trusted source list: Ordered by level of trust	This list contains all of the sources applied to the selected table or field. Prioritize the sources using the up and down arrows. The source at the top has the highest priority.
3) Trusted source list: Available trusted sources	This list contains all sources available to apply to the selected table or field. Use the right and left arrows to move sources between the lists.

1. Update the matching table configuration to use the correct source field:


- a. Open the desired table configuration by selecting the **Preview** button as shown below from the trusted sources screen, or do so from **Table activation and settings**.

Note:

Before navigating away from this screen save your changes.



- b. On the **Main** tab, select the field indicating the source from the **Trusted source field** menu:



- c. Save and close the record.

Now when you configure merge policies for this table and choose the **Most trusted source** option, the add-on will prioritize survived values according to the trusted source list you just created. Note that when selecting the values to survive, the add-on prioritizes those configured at the field level first and the table level second.

Data steward decisions as sources

You can enable user decisions to override the order of trusted sources for future merge operations by adding a **Data Steward decision** to a trusted source list. When a user runs the [Supersede survivorship](#) service they can override a field merged as the result of adhering to a trusted source list. How the add-on handles the same field in future merges depends on the position of the **Data Steward decision** source in the trust list.

Attention:

If a merge policy uses the **Customize source value for new golden** option to use a specific value, users cannot change source field values. This would create a direct conflict with the merge policy. If you want to allow users to override a merge decision, set the **Customize source value for new golden** to .

When a user overrides a merge decision and the **Data steward decision**:

- is at the top of the list, the add-on "learns" and uses the value obtained from the same source the next time that record is involved in a merge operation.
- is *not* at the top of the list, the add-on will use values from sources with a higher trust level over the one chosen by the user.
- has trusted sources below it in the list, the add-on uses the value from the user decision in the future over those below it.

The following example includes two fields with values from different sources. Notice the result when the add-on applies different trusted source lists to an auto merge performed *after* a user executes the **Supersede survivorship** service:


		Color:	Number:	Source:
	Record 1	Red	One	CRM
	Record 2	Green	Two	ERP
		Merge results		
If the trusted source order is set to: 1) CRM 2) Data Steward decision 3) ERP		Initial auto merge	Data steward decision	Next auto merge
		Red	Green	Red
		One	One	One
If the trusted source order is set to: 1) Data Steward decision 2) CRM 3) ERP		Initial auto merge	Data steward decision	Next auto merge
		Red	Green	Green
		One	One	One

Matching with relationships


The add-on allows you to perform matching operations on fields that derive their value from related tables through direct foreign key relationships and join tables. You can configure these fields in the decision tree the same as other fields, with one prerequisite: you must create a configuration for the field on a matching policy's **Matching field** tab. The field is not available to select in a decision tree node until you complete this step.

The following steps outline how to configure this option:

1. Create a matching field:
 - a. Select the **Administration** tab and navigate to *TIBCO EBX® Match and Merge Add-on > Table activation and settings*.
 - b. Open the desired table and select its **Matching policies** tab.




c. Open the desired matching policy and select the **Matching fields** tab.



d. Create a new record:

The image below and following table highlight some of the options:



Use relationships

Choose **By foreign key** or **Over a link table**.

Note:

When selecting the **By foreign key** option, the you can only specify a single hop foreign key. In other words, it must point to a regular field and not another foreign key.

Location

Select the option that corresponds to the related table's location. Note that the current release only supports the **Same dataset** option.

Matching table foreign key and Related table foreign key

Matching table foreign key: Select the wand icon and choose the foreign key you want to match through. When matching using a link table, this is the relationship between the table where matching is configured and the link table.

Related table foreign key: This option is only available when the **Use relationships**


	property is set to Over a link table . Select the wand icon and choose the relationship between the link table and table containing the source value.
Field	Select the field in the related table where you want to apply matching.
Exclude records where	Specifies values to exclude records from matching operations. See Excluding records for more information.
Pre-processing	Some parts of the matching process are handled by the core EBX® engine and others by the add-on. Pre-processing refers to the processing done by EBX® to find similar records before the add-on matches them. See Understanding matching operation processing for more information. Search strategy and Weight: Select the search strategy to use and provide a weight. Available search strategies are defined at the data model level. The weight sets the level of importance placed on this field as compared with others. The higher the value, the more weight the field has.
Null value management	Use the options in this group to determine how null values are handled during a matching operation. Attention: When enabling options to handle null values, you should set the field's Weight to a lower value relative to other fields. Not doing so can result in unexpected matching results.

1. Open the desired decision tree and double-click to edit a data comparison node.
2. Configure the node to use the foreign key field by adding it to the node's **Fields**:
 - Set **Use relationships** to **By foreign key** or **Over a link table**.
 - Use the **Field** menu to select the foreign key field configured in the **Matching fields** tab.

Attention:

At least one of the fields used in the decision tree must be located in the table where matching is configured and active.

- Define the **Algorithm** and **Weight** options as needed.



3. Save and close the **Data comparison node** configuration.
4. Save and close on the **Decision tree** tab.

Understanding matching operation processing

- Overview
- Pre-processing and cluster generation
 - Purpose and principles

- How similarity is determined
- Defining criteria and exclusions
- Balancing field weights
- Example: optimizing similarity criteria
- Adjusting parameter settings
- Pre-processing optimizations reduce cluster size


Overview

The add-on and the EBX® core share responsibility for matching process execution. The core product pre-processes the list of records to match to find similar records and passes them to the add-on. EBX® finds similar records by taking the base record, or record to match against other records, and applying certain parameter settings.

If a field configured to participate in matching meets either of the following criteria, it is automatically disabled during the pre-processing phase to improve performance:

- The field's type is `boolean`.
- The entire column is `null`.

When input consists of a selection of records, the output is a *group* of similar records. If the input is an entire table, additional processing is performed and the output is *pairs* of similar records. The add-on processes the output from EBX® using the logic defined in a matching policy's decision tree to determine whether values match. The diagram below visualizes this process:



Pre-processing and cluster generation

To ensure high performance and scalability, not every record is compared against all other records. Instead, a pre-processing phase (clustering) creates candidate clusters for matching.

Purpose and principles

Clustering acts as a coarse filter. The goal is to reduce how often the add-on's decision tree logic runs. It does this by grouping only likely matches into clusters and filtering out obviously unrelated records. The logic is also tuned to keep false negatives (missed matches) to a minimum.

How similarity is determined

During pre-processing, EBX® iterates over the records to match and uses a full-text search to find similar records. Generally, every matching field configured in the matching policy acts as a criterion for similarity, where the system breaks down field values into searchable terms (or keywords). To be considered similar and compared within the decision tree, a record pair typically needs to satisfy a minimum threshold—for example, sharing at least 1 significant term out of 3.

Defining criteria and exclusions

Not all fields make good similarity criteria. Fields with low cardinality (few unique values) or very common values can create massive, inefficient clusters. You must manually exclude fields that are not unique enough to identify candidates (e.g., a "Country" field where 90% of customers are in "US", or a "Gender" field). To manually exclude a field, set the Weight to 0 in the matching policy's **Matching fields** tab. This ensures the field is used for scoring in the decision tree but not for picking candidate records during clustering.

The system also applies automatic exclusion. Fields defined as Boolean (e.g., IsActive) are automatically disabled during the pre-processing phase to improve performance. If an entire column is null, it is also automatically excluded.

Finally, specific common words (stopwords) must be filtered out to prevent irrelevant matches. Without this, records sharing generic terms like "Inc." or "LLC" might be identified as similar even if they are completely unrelated. This also applies to single-letter keywords often found in composed names (e.g. "George R.R. Martin"), unless filtered, the letter "R" could trigger matches with any other name containing that initial.

Balancing field weights

By default, all matching fields are assigned a weight of 1. Keep the default value unless you notice quality issues in your matching results. Only adjust weights if specific fields need to be prioritized to improve the relevance of the clusters. As mentioned above, set the weight to 0 only when you want to exclude a field from pre-processing.

When adjustments are necessary, assigning appropriate weights becomes crucial for the coarse filter to work effectively. You should only give higher weights to fields that are highly distinctive (like a Tax ID or Email) and lower weights to fields that are less specific or prone to errors (like First Name).

Example: optimizing similarity criteria

Consider a matching policy for Corporate Customers that uses the following fields: Company Name, City, Region, and IsActive.

FIELD	EXAMPLE VALUE	CLUSTERING BEHAVIOR	CONFIGURATION NOTE
IsActive	True	Ignored	Automatic, no configuration required. Boolean fields are never used to find similar records.
Region	North-East	Ignored	Manual configuration is required. Because "North-East" is too generic and would cluster too many unrelated records, an administrator must set the Weight of the matching field to 0.
City	Chicago	Used	Used as a similarity criterion. Records sharing "Chicago" are candidates.
Company name	Acme Inc.	Used, but optimized	Used as a similarity criterion. Records sharing "Acme" are candidates. Stopword: The term "Inc." is configured as a stopwords and ignored to prevent matching with each occurrence of "Inc." in the column.

In the above example, records are picked as "similar" candidates only if they share a meaningful term (e.g. "Chicago" or "Acme") in the City or Company Name fields. System resources are not wasted comparing records just because they are active (True), or in the "North-East".

Performance tip: Selecting the right fields for clustering is critical. For detailed instructions on tuning these settings to avoid bottlenecks, refer to the [Performance recommendations](#) section to optimize performance.

Adjusting parameter settings

Adjusting parameter settings changes the criteria used by EBX® to find similar records. As a result, the output records passed to the add-on for processing can change. Fine-tune the settings to obtain desired results. However, keep in mind that it is wise to find a balance between the trade-off of restrictive settings vs nonrestrictive. Criteria that is too strict can result in not passing records to the add-on that a decision tree would consider a match.

To edit parameter settings:

1. Access the configuration settings for the matching field you want to update:
 - a. Navigate to *Administration > TIBCO EBX® Match and Merge Add-on > Table activation and settings*.
 - b. Open the table configuration and matching policy containing the field.
 - c. Select the **Matching fields** tab and open the field.
2. Use the following options under **Pre-processing** to change parameter settings:
 - a. **Search strategy:** Sets the search strategy used to group similar records during the pre-clustering process. The list only includes search strategies defined in the data model containing the selected field. See the *EBX® Reference Manual* for more information about search strategies.
 - b. **Weight:** Specifies the weight assigned to this field during pre-processing by EBX®. When a matching operation includes multiple fields, the weight sets the relative importance of each field's score in determining whether records are similar. Note that this weight does not affect the weighted average in the decision tree comparison nodes.

See the information above on [balancing field weights](#).
 - c. **Null value management:** Determines the behavior when comparing two records and one or both of them has a Null value in this field. When enabling matching of null values, adjust the weight for the field so that it is lower as compared to other fields. Otherwise, it might not get included during pre-processing with other similar records. Using the **Do nothing** option when your dataset contains many null values can help performance and result relevancy. This option ignores null values and prevents records that are not similar, except for their

null values, from being considered similar. Essentially, it helps prevent false positives from passing to the decision tree for unnecessary evaluation.

3. Save and close to keep your changes.

Pre-processing optimizations reduce cluster size

Where possible the pre-processing phase automatically reduces the cluster size to optimize performance. During this phase, the decision tree is analyzed as follows:

- The system identifies any fields determined as "mandatory". These are fields whose similarity is required to consider records as matches or suspects. Any records without similarity in these fields are excluded from further processing, which can significantly improve performance by reducing the number of candidate records. For example, the Customer_Email field is considered as mandatory in a decision tree where the field is included in all paths that lead to a MATCH or SUSPECT output. In this case, records that do not have the same email address will never be considered as candidates, and are not passed to the decision tree for evaluation.

Currently, this behavior only applies to data comparison nodes that use the **All fields match** evaluation and the following algorithms: **Exact**, or **Full text** with a minimum score of 100%.

- The system analyses each decision tree path that leads to a MATCH or SUSPECT output to identify groups of fields where similarity is essential for records to match. A minimum required score is computed from the smallest group of fields. Records must meet, or exceed this score. Otherwise, they are not passed to the decision tree for evaluation. Take for instance, a decision tree with five fields configured with the same weight, and at least three of the fields must be similar for records to match. In this case, the minimum similarity score is 60% to pass candidate records to the decision tree.

Perspectives and workflows

EBX Match and Merge Add-on services can be applied to EBX® perspectives and workflows. For information on creating perspectives and workflows, see the EBX® product documentation. The add-on also allows you to specify that a workflow launches automatically when inline matching is triggered. See [Creating a matching policy](#) for instructions on enabling the option.

Attention:

When configuring a workflow some of the parameters require specific values. Be sure to check each parameter's tooltip to ensure expected behavior.

The following table provides descriptions of the add-on services you can use in perspectives and workflows:

SERVICE	DESCRIPTION
[ebx-addon-mame] Merge (workflow only)	Opens the add-on's Merge view to enable users to manually merge records.
[ebx-addon-mame] Run match (workflow only)	Executes the Run match service on a set of records or entire table. Note: This service is available as a workflow User task and Script task.
[ebx-addon-mame] Inline matching (workflow only)	Executes matching during the creation or update process. Users must review any potential duplicate records before continuing.
[ebx-addon-mame] Wait for processes to finish (workflow only)	Waits to execute the next workflow step until current processes are finished.
[ebx-addon-mame] Manage group (workflow only)	Opens the Manage group screen where users can view and manage results related to individual records.
[ebx-addon-mame] Consolidated view (for perspectives and workflows)	Opens the Consolidated view where users can view matching results and access other services to manage results.
[ebx-addon-mame] Search before create (for perspectives and workflows)	Executes the Search before create service to allow users to search for potential duplicate records prior to creating them.

SERVICE	DESCRIPTION
[ebx-addon-mame] View lineage (for perspectives and workflows)	Opens the Lineage tab which shows the data lineage that resulted in the creation of the golden record.


Displaying matching metadata

Matching metadata is automatically included in tables. However, you can create a view that displays add-on metadata in a custom location. This might be useful for larger tables to prevent users from having to scroll through many columns to locate the metadata. They could then quickly sort records by state to see if any action is required and access record groups. Metadata applied by the add-on includes:


- State
- Group
- Auto-created

To add matching metadata to a table view:

1. Open the table that you want to display matching metadata.
2. Open the **View** menu and select **Create a new view**.




- Optionally, provide a name and select **Next**.
- The three matching metadata categories are automatically included in the **Displayed columns** list.



- Make any additional desired changes to the columns displayed and their order. Select **Create and apply**.

The new columns displays in the specified location as shown below:



Matching algorithms

- Included algorithms
- Search strategies and recommended algorithms
- Sample match scores by algorithm

Included algorithms

The add-on includes several predefined algorithms to use in matching operations. Each algorithm has advantages and disadvantages depending on the type of data being searched. The following table provides descriptions for the algorithms included with the add-on and any input parameters they might have:

ALGORITHM	DESCRIPTION
Hybrid fuzzy	This algorithm is designed to match two strings by evaluating their similarity using seven different criteria. Each criterion is assigned a weight value between 0.0 and 1.0. A criterion is also assigned a priority

ALGORITHM	DESCRIPTION
	order that is applied during the comparison process. When comparing two strings, the algorithm will go through all seven criteria in the priority order. If the whole string or parts of it meet the condition of a particular criterion, the weight assigned to that criterion will be used in the final formula. If conditions are not met, that criterion is skipped, and the algorithm moves to the next one. Once all the criteria have been evaluated, the final similarity score is calculated based on the average weight of the remaining criteria. If the resulting score is greater than the threshold value defined by an administrator in the decision tree's comparison node, the two strings will be considered a match. Points of note about this algorithm: • Does not return a 100% matching score on synonyms. • Strings that are similar, but contain characters like spaces or dashes do not return a 0% matching score. • When strings fall within the acceptable distance score they are not considered as the exact same. For example, a distance of 1 returns .95, a distance of 2 returns .9, etc.
Phonetic full text	An algorithm best used for strings. It can recognize the phonetic equivalent of two

ALGORITHM	DESCRIPTION
	words spelled differently using the Beider-morse phonetic tokens. The comparator takes into account the synonyms and stop words defined in the data model.
Beider-morse	A phonetic algorithm for short strings (e.g. proper names) that is able to recognize a phonetically equivalent of two words written in a different way. The "sounds-alike" test is based not only on the spelling but on linguistic properties of various languages.
Full text	An algorithm best used for strings. It can find the case-insensitive exact matches of the words in the compared values. The comparator takes into account the synonyms and stop words defined in the data model. Note: Values that contain only stop words are not supported. Additionally, this algorithm is only compatible with String and text data types.
Fuzzy full text	An algorithm best used for strings. It can find the similar and fuzzy matches of the words in the compared values and is based on the Levenshtein distance algorithm. The comparator takes into account the synonyms and stop words defined in the data model. Note: Values that contain only stop words are not supported. Additionally, this algorithm is only compatible with String and text data types.

ALGORITHM	DESCRIPTION
Exact	An algorithm to match data that should be exact (e.g. code). This algorithm returns a matching score of 100% for values that are exactly the same. Matching using the Exact algorithm is case-sensitive by default for fields with String and Text data types. You can set it to case-insensitive when configuring data comparison nodes in the decision tree.
JaroWinkler	A distance algorithm for short strings (e.g. proper names) that tallies the number of characters in common and places a higher emphasis on differences at the start of the string. Input parameters: The 'Threshold' parameter values should be from 0.0 to 1.0. The parameter determines when a Winkler bonus should be added. Decreasing this parameter might result in an increased score. For example: With the keyword 'Fra' and data 'France', if the 'Threshold' = 0.7 score = 88.33. If the 'Threshold' = 0.9 score = 83.33. Default value: 0.7
Levenshtein	A distance algorithm for short strings that works well when only few differences between the values are expected. For example, this works well for dialects spoken in a particular part of the country, or by a specific group of people.
Soundex	A phonetic algorithm for short strings (e.g. proper names) that indexes the strings based on the way they sound in English

ALGORITHM	DESCRIPTION
	rather than the way they are spelled. The homophones are encoded to the same representation so that they can be matched despite minor differences in spellings.
Range	An algorithm to match values within a predefined range. The two values are considered a match if the distance between two values is within the range. Note: This algorithm is only available for numeric and date/time fields.

Search strategies and recommended algorithms

Some matching algorithms are not compatible with certain search strategies and could return inaccurate matching results. Thus, when configuring a comparison node in a matching policy's decision tree, the list of available algorithms is filtered depending on a field's configured search strategy. This helps prevent an incompatible configuration. Additionally, the combination of a field's configured search strategy in the Data Model Assistant and its matching algorithm can impact performance. See [Performance recommendations](#) for more information.

The following table shows the compatibility between search strategies and algorithms:

SEARCH STRATEGY	ALGORITHMS
Default search template (set in DMA)	
Default full text search strategy or any other configured for the default template.	• Exact • Levenshtein • JaroWinkler • Full text • Fuzzy full text

SEARCH STRATEGY	ALGORITHMS
	Not recommended but still available: • Beider-morse • Soundex • Phonetic full text
Non-default search templates	
Tese Levenshtein	• Exact • Levenshtein • JaroWinkler • Full text • Fuzzy full text
Tese NGram	• Exact • Levenshtein • JaroWinkler • Full text • Fuzzy full text
Tese Soundex	• Exact • Soundex • Beider-morse • Phonetic full text
Tese Double Metaphone	• Exact • Soundex • Beider-morse • Phonetic full text

SEARCH STRATEGY	ALGORITHMS
Jaro-winkler	• Exact • Levenshtein • Jaro-winkler • Full text • Fuzzy full text

Sample match scores by algorithm

The following use cases include sample input values and each algorithm's score result after comparing the input:

USE CASE	INPUT VALUES	SCORE/ALGORITHM
Single typo	1. Robert Downey 2. Robetr Downey	Beider-morse: 1 FuzzyFullText: 1 PhoneticFullText: 1 Jarowinkler: 0.98 Levenshtein: 0.85 HybridFuzzy: 0.75 FullText: 0.50 Soundex: 0.50 Exact: 0
Double typo	1. Robert Downey 2. Robetr Downye	FuzzyFullText: 1 Jarowinkler: 0.96 Levenshtein: 0.69

USE CASE	INPUT VALUES	SCORE/ALGORITHM
		HybridFuzzy: 0.50 PhoneticFullText: 0.50 Beider-morse: 0.50 Soundex: 0.50 FullText: 0 Exact: 0
Extra strings	1. Robert Downey 2. Robert Downey Jr	FuzzyFullText: 1 FullText: 1 HybridFuzzy: 1 PhoneticFullText 1 Jarowinkler: 0.99 Levenshtein: 0.81 Beider-morse: 0.80 Soundex: 0.80 Exact: 0
Stopwords (Mr-Jr)	1. Robert Downey 2. Mr Robert Downey Jr	FuzzyFullText: 1 FullText: 1 HybridFuzzy: 1 PhoneticFullText: 1 Jarowinkler: 0.82 Levenshtein: 0.68 Beider-morse: 0.67

USE CASE	INPUT VALUES	SCORE/ALGORITHM
		Soundex: 0.67 Exact: 0
Mixed word order	1. Jr Robert Downey 2. Jr Downey Robert	FuzzyFullText: 1 FullText: 1 PhoneticFullText: 1 Beider-morse: 1 Soundex: 1 HybridFuzzy: 0.95 Jarowinkler: 0.92 Levenshtein: 0.38 Exact: 0
Soundex words	1. Hayley McDonnell 2. Hailee McDonald	Soundex: 1 Jarowinkler: 0.83 Levenshtein: 0.69 FuzzyFullText: 0.50 PhoneticFullText: 0.50 Beider-morse: 0.50 HybridFuzzy: 0.25 FullText: 0 Exact: 0
Phonetic words	1. John has four apples 2. Jon haz for appls	PhoneticFullText: 1 Beider-morse: 1

USE CASE	INPUT VALUES	SCORE/ALGORITHM
		Soundex: 1 FuzzyFullText: 1 Jarowinkler: 0.81 Levenshtein: 0.80 HybridFuzzy: 0.50 FullText: 0 Exact: 0
Synonyms	1. Mike 2. Michael	PhoneticFullText: 1 FuzzyFullText: 1 FullText: 1 Jarowinkler: 0.78 HybridFuzzy: 0.50 Levenshtein: 0.43 Beider-morse: 0 Soundex: 0 Exact: 0
Special characters	1. Alex: Bin? Alex!" 2. ----Bin Alex; Bin@	FuzzyFullText: 1 FullText: 1 PhoneticFullText: 1 Beider-morse: 1 Soundex: 1 HybridFuzzy: 0.63

USE CASE	INPUT VALUES	SCORE/ALGORITHM
		Jarowinkler: 0.53 Levenshtein: 0.28 Exact: 0
Range	Use case 1 :- "01/02/2021 23:00:00", "01/02/2021 22:00:00" (Range 1 day) Use case 2 :- "01/02/2021 23:00:00" , "02/02/2021 01:00:00" (Range 1 day) Use case 3 :- ""01/02/2021 23:00:00", "01/02/2021 23:00:03" (Range 2 sec) Use case 4 :- "10", "16" (Range 4)	Use case 1 :- 1 Use case 2 :- 0 Use case 3 :- 0 Use case 4 :- 0

Performance recommendations

- [Overview](#)
- [EBX® optimization](#)
 - [Identifying indexing issues](#)
- [EBX Match and Merge Add-on optimization suggestions](#)
- [Troubleshooting](#)

Overview

The EBX Match and Merge Add-on and EBX® core product work together to handle matching tasks. Due to this shared responsibility, the biggest performance gain comes from making optimizations in both the add-on *and* the core product. For more details on the roles of each during the matching process, see [Understanding matching operation processing](#).

The following topics describe:

- [EBX® optimization](#)
- [EBX Match and Merge Add-on optimization suggestions](#)
- [Troubleshooting](#)

EBX® optimization

To ensure the best performance when matching, it is important to follow EBX® performance guidelines, especially for high volumes of data. It is of particular importance to allocate less than 1/3 of the application server's available memory to the JVM. This leaves the rest free for file system caching. Additionally, the pre-processing phase is optimized to benefit from multiple CPUs. This means that CPU oriented hardware can improve matching performance. See *EBX® documentation > Administration Guide > Installation and configuration > Configuration notes* for the most up-to-date details.

Identifying indexing issues

The search strategies applied to fields involved in matching can have a large impact on performance. The following topics present some best practices that can help prevent performance issues:

- [Stop words](#)
- [Prefer code search strategy](#)
- [Avoid empty or null value placeholders](#)
- [Use fuzzy and phonetic search strategies sparingly](#)
- [Child dataspace](#)

Stop words

String values that contain common tokens greatly increase the number of irrelevant match candidates found when pre-processing and slow the matching process. To alleviate this, you can [define stop words](#), which EBX® ignores during pre-processing (i.e. they are not indexed). Stop words are registered, then defined in the search strategy of an `xs:string` field in the **Search** extension of the Data Model Assistant.

The best candidate terms for stop words are frequently found values. For example, 'street' and 'blvd' make good stop words because they exist in almost all records related to addresses. The

Troubleshooting section describes how to get the most frequent values for a field configured to use matching.

Improper selection of stop words can negatively impact performance, particularly when too few values are left in a record to establish similarity. If a value is composed solely of stop words, it indicates that the words have significant semantic meaning and should not be considered stop word, despite their frequency.

Note:

The result of a matching process is not affected by a set of well-chosen stop words. They are ignored by the core product during pre-processing, but are still taken into account when executing the add-on's decision tree logic.

Prefer code search strategy

The **Name** and **Text** search strategies tokenize values to support full-text search. Nevertheless, some data are ill-suited for tokenization. This can lead to an abundance of common tokens, which slows the matching process. This especially holds true for `xs:string` fields holding identifiers, dates, phone numbers or emails. For these cases, the preferred method is to use the **Code** search strategy, which ensures that the whole value is indexed as a single token.

Note:

By default, the **Code** search strategy uses the *contains* operator when searching. This is expensive in terms of performance on large volumes. Instead, update the strategy to use *starts with* when matching mid-sized to large volumes.

Avoid empty or null value placeholders

When a value is null or empty, do not use a placeholder like 'Not available', 'N/A', or default values like '1', 'xxx', etc. Those placeholders are common tokens, and slow down matching. It is recommended to leave the value empty instead.

Use fuzzy and phonetic search strategies sparingly

The fuzzy and phonetic search strategies are slower than the default, built-in strategies. In the context of matching, this impacts:

- Core product performance when gathering similar records during pre-processing. The cost is mainly due to the increased number of results for a fuzzy strategy even if only applied to one decision tree node.
- Add-on performance is impacted as fuzziness can add a lot of false positives in the results. The increase in the number of records passed to the add-on can slow the process. Since these additional records might not meet conditions of other comparison nodes, the add-on is required to process additional records with no guarantee of finding matches.

If the conditions on other comparison nodes are met, then a record is considered 'similar' during pre-processing, even without applying a fuzzy strategy to a field. Hence, we recommend keeping the default search strategy when using fuzzy or phonetic algorithms in the decision tree, and when the decision tree contains at least 4 fields.

However, if your matching policy still requires a fuzzy or phonetic search strategy, the following information can help to mitigate the performance impact:

- Jaro-Winkler is the most expensive.
- The Levenshtein search strategy accepts parameters to tune its behavior, which impact performance. First, Levenshtein is an edit-distance based algorithm. If distance 2 is too expensive, try decreasing the distance to 1. Second, there is a parameter to specify whether the values are tokenized or not. Without tokenization, the algorithm behaves like a fuzzy **Code** strategy, and is less expensive. With tokenization, it behaves like a fuzzy **Text** strategy. The fuzziness is applied to all tokens, but it greatly increases the number of results returned and is more expensive. The parameters are configured in the **Search** extension of the DMA
- The phonetic parameter of **Name** and **Text** search strategies uses the Beider-Morse phonetic algorithm. This is the least expensive.

Child dataspace

If you have multiple child dataspace, be sure that you do not run matching operations on these simultaneously. It can lead to poor performance and inaccurate results.

Transaction management

In scenarios with successive modifications, it is advised to batch these modifications into a single transaction. Batching transactions helps minimize the overhead associated with multiple transactional operations. This practice is particularly beneficial and can ensure optimal performance and resource utilization in environments where numerous modifications occur in succession.

EBX Match and Merge Add-on optimization suggestions

The following matching policy recommendations can help improve performance:

- **Decision tree:** Before performing matching on an entire table with a large volume of records, test the matching policy on a subset of records (less than 50k records). And:
 - Check that the number of matches and suspects is consistent with your dataset. If it is not consistent, start the troubleshooting process with the decision tree. Begin by reviewing the nodes leading to the **Match** or **Suspect** output.
 - When a grouping or merging phase is slower than usual, it is likely the symptom of other configuration issues.
- **Weight:** By default, all matching fields are assigned a weight of 1. This value should only be increased in rare cases, such as the following:
 - If the first comparison node in a decision tree includes a field that strongly influences the outcome of a match operation, increasing its weight can improve performance and reduce false negatives.
 - If you see false negative (records that should be considered as matching or suspect but are not), you can experiment with increasing the weights of fields used in the initial nodes of the decision tree.

A more common use case is to use weights to exclude records during pre-processing. The pre-processing phase estimates the similarity between records based on the value of their matching fields. Pre-processing is sensitive to frequent tokens and values, as it increases the number of irrelevant match candidates. To improve matching performance, set the weight to 0 for fields that have a limited number of options. For example, fields that are boolean, enumerations, or flags. A field with a weight of 0 is not considered when estimating similarity during the pre-processing phase, but it is still taken into account when executing the decision tree comparisons.

- **Exclude records:** By defining a set of criteria that excludes certain records from the matching operation, you can eliminate a subset of records that the add-on must process. This can lead to performance improvements. See [Excluding records](#) for more details.
- **Matching on related tables and business objects:** Matching configurations that include relationships between multiple tables, or configurations applied to business objects introduce complexity that significantly impacts performance. For mid to large-sized datasets, it's advisable to denormalize the data model, enabling a streamlined single-table matching policy. This simplifies the comparison logic and reduces the overhead associated with relationship-based queries, enhancing the overall speed and efficiency of matching operations.

Troubleshooting

Activating the logs can help you investigate matching process performance. The kernel log contains key indicators for each field configured to use matching and offers valuable clues for applying performance recommendations. For example:

- Number of occurrences of the most frequent token: If the count is significantly high compared to the table size, it may indicate the absence of essential stop words. When the property `ebx.log4j.category.log.kernel` is set to `DEBUG`, the 10 most frequent tokens are logged, and should be considered as stop words. See [Stop words](#) for more information.
- Number of occurrences of the most frequent value: when the field contains a few distinct values with a high frequency, the field is not suitable to estimate the similarity between records. Consider ignoring such fields during preprocessing. See the bullet point discussing **Weight** in [EBX Match and Merge Add-on optimization suggestions](#).
- Number of null values. Excessive null values in a field, especially when the matching field is configured to match with other null values or any value, can significantly impact performance. A record with null values will appear similar to many others. This leads to a lot of decision tree executions. In this case, it is better to configure null values to **no match**.
- The **Code** search strategy with the *contains* parameter: The *contains* operator is expensive on large volumes. Modify it to *starts with* when matching mid to large volumes.

The log also provides precise information about matching progress: number of records, average throughput, percentage of completion. This is useful to estimate the duration of the matching phase. It can also help you determine whether pre-processing is slower than it should be, considering the volume of data.

Replicating matching metadata

- [Overview](#)
- [Configuring metadata replication](#)

Overview

Replicating your matching metadata allows you to use third-party tools to query the metadata. For example, you might want to expose matching metadata to a BI dashboard for analysis.

When enabling replication, you must configure settings at the table level for individual tables and business objects. You specify which table publication to get metadata from, the method used to trigger table refresh, and the table name in the database for the replicated metadata table.

Some of the information included in the replicated metadata is each record's:

- Group id
- State
- Golden id (For merged records)
- Merging process id

Note that the State only includes the identifiers used for matching states. The ids are mapped to their state in the following table:

ID	STATE
0	UNSET
1	GOLDEN
2	MERGED
3	PENDING APPROVAL
4	DELETED

Configuring metadata replication

To enable matching metadata replication:

1. Add replication to a matching table configuration:

- a. Open the desired matching table configuration:
 - i. For individual tables: Navigate to *Administration > TIBCO EBX® Match and Merge Add-on > Table activation and settings* and open the desired matching configuration.
 - ii. For business objects: Navigate to *Administration > TIBCO EBX® Match and Merge Add-on > Business object activation and settings*, open the desired configuration, and select the table you want to replicate.
 - b. Select the **Replications** tab.
 - c. Fill in the required information. Note that the **Table name in database** property must specify a unique name in the database. Additionally, this name must conform to any database requirements.
 - If set to **On commit**, the replication metadata is pushed to the database after matching executes and commits happen in the EBX® repository.
 - If set to **On demand**, a user must manually push the replication metadata to the database from the dataset's **Actions** menu by selecting: *Replication > Refresh replicas*.
2. Publish your data model. If you come back to this replication configuration and make changes, you must re-publish your data model to apply them.

Attention:

When setting up replications, the maximum length for table's primary key must be specified (not left undefined) in the Data Model Assistant (DMA).

Scheduling matching metadata purges

Overview

To prevent performance degradation, you can remove outdated matching metadata. Business users can perform this action manually using the **Purge old result** service. See [Purging old results](#) for more information. As an administrator you can automate this service using the EBX® scheduler.

When you use the scheduler, you'll create:

- a *task* for each table whose matching metadata you want to purge.
- a *schedule*, or time increment that defines how often the task runs.

- a *scheduled task* that links the task to the schedule and enables or disables execution.

To schedule regular purges of matching metadata:

1. Navigate to *Administration > Technical configuration > Task scheduler > Tasks* and create a new task.
2. Provide a name, optional description, and in the **Specification** group:
 - For **Module** select **Match and Merge**.
 - For **Class name**, enter:


```
com.orchestranetworks.addon.mame.scheduler.SchedulerPurge
```

3. Save, but do not close the task.

After successful save, the **Specification** group displays the **Dataspace**, **Dataset**, and **Table** fields.

4. Use the fields displayed to specify the location of the table where you want the **Purge old results** service to run. Note that you must enter the full table path, for example:

```
/root/<your_table_name>
```

.
5. Save and close the task.
6. Go to **Schedules** and create a new schedule.
7. Add a name and use the **Cron expression** field to write a cron expression that defines the schedule increment. Run the purge weekly or every two weeks for most workloads. See the *TIBCO EBX®* documentation for more information about *Cron expressions* and their syntax.
8. Save and close the schedule.
9. Go to **Scheduled tasks** and create a new record:
 - Select the **Task** created in earlier steps.
 - Select the **Schedule** increment created in the previous step.
 - Specify whether the scheduled task is active.
10. Save and optionally close the record.

When the task is active, the UI displays the next planned execution time.

Developer Guide

- [REST overview](#)
- [Java API](#)

REST overview

The TIBCO EBX® Match and Merge Add-on allows access to its services using a RESTful API. The add-on follows the OpenAPI Specification (OAS) to document its REST API. This allows you to use tools such as Swagger to generate and interact with our REST documentation. It is not our intent to document how to use these tools here. However, the following example shows one of the more basic ways to view the REST API documentation.

To generate the add-on's OAS:

1. Use a tool such as Postman to send a GET request to the following endpoint on your EBX® server: <http://:/ebx-addon-mame/rest/v1>.
2. Ensure the response is formatted in JSON.
3. Save the response to a file.
4. Open a web browser and navigate to Swagger Editor: <https://editor.swagger.io/>.
5. Select *File > Import file* to import the add-on's specification.

The editor displays the formatted text on the left and HTML on the right.

```

1 openapi: 3.0.1
2 info:
3   title: EBX Match & Merge Public API
4   description: >-
5     EBX Match & Merge Add-on allows to leverage EBX native abilities by
6     supporting the data quality dimension with data matching and data merging.
7     Data is available in JSON format. This swagger cover every operations
8     available from the business user UI of EBX Match & Merge add-on
9   contact:
10    email: support@tibco.com
11    version: '3.0'
12   paths:
13     /rest/v1/groups/{groupId}:
14       get:
15         tags:
16           - group
17         summary: Returns all matches and suspects of a group.
18         description: >-
19           Returns all matches and suspects of a group, in order to have detailed
20           information of this group.
21         operationId: getGroupById
22         parameters:
23           - name: groupId
24             in: query
25             schema:
26               type: string
27           - name: dataset
28             in: query
29             schema:
30               type: string
31           - name: tablePath
32             in: query
33             schema:
34               type: string
35           - name: groupId
36             in: path
37             required: true
38             schema:
39               type: number
40         responses:
41           '200':
42             description: detail of a group of matched records
43             content:
44               application/json:
45                 schema:
46                   $ref: '#/components/schemas/GroupDetail'
47     /rest/v1/groups:
48       get:
49         tags:
50           - groups
51         summary: Returns the list of groups of matched records that need action
52         description: >-
53           Returns the list of groups that needed action from a data steward, as to
54           know groups that need to be merged, has suspects or contains false
55           positives. The threshold of groups return here is limited to 50 per
56           category and it's not possible to override it for the client
57         operationId: getGroups
58         parameters:
59           - name: dataset
60             in: query
61             schema:

```

EBX Match & Merge Public API ^{1.0} OAS3

EBX Match & Merge Add-on allows to leverage EBX native abilities by supporting the data quality dimension with data matching and data merging. Data is available in JSON format. This swagger cover every operations available from the business user UI of EBX Match & Merge add-on

[Contact the developer](#)

group

GET /rest/v1/groups/{groupId} Returns all matches and suspects of a group.

groups

GET /rest/v1/groups Returns the list of groups of matched records that need action

matchingProcess

GET /rest/v1/matchingProcesses/{matchingProcessId} Returns information of a matching process

DELETE /rest/v1/matchingProcesses/{matchingProcessId} Cancel a matching process execution.

matchingProcesses

POST /rest/v1/matchingProcesses run match operation on a whole EBX table or a selection of records

POST /rest/v1/matchingProcesses/search searches before create any record into an EBX table

POST /rest/v1/matchingProcesses/simulate Simulate a matching execution for an existing record into an EBX table.

mergingProcesses

POST /rest/v1/mergingProcesses Merge a group

DELETE /rest/v1/mergingProcesses/{dataset}/{pathSegments} Unmerge a record

Java API Reference

- TIBCO EBX® Add-ons 6.2.3 Java API

User Guide

- [Matching and merging overview](#)
- [Matching and Merging Walkthrough](#)
 - [Overview](#)
 - [Part 1: Initiate matching](#)
 - [Part 2: Using the Consolidated view](#)
 - [Part 3: Merging suspects](#)
- [General Tasks and Views](#)
 - [Automatic Matching](#)
 - [Evaluating matching](#)
 - [Searching before taking action](#)
 - [Manual Matching](#)
 - [Purging old results](#)
 - [Managing groups](#)
 - [Merging Records](#)
 - [Merging business object records](#)
 - [Reversing a merge operation](#)
 - [Consolidated view](#)
 - [Changing record state](#)
 - [Viewing add-on metadata](#)

Matching and merging overview

- [Overview](#)
- [Matching](#)
- [Merging](#)

Overview

The TIBCO EBX® Match and Merge Add-on's purpose is to locate and merge duplicate data into a *golden record*. A golden record represents the most truthful, accurate, and up-to-date data. As a business user, the way you interact with the add-on to match and merge data depends on configuration settings.

Matching

The add-on's *matching operations* entail looking at data values to determine whether duplicate records exist. Administrators will configure how the add-on accomplishes this. For configuration instructions, administrators can use the [Quick Start Guide](#).

At any time, you can manually initiate a matching operation. When administrators enable inline matching, the add-on automatically matches newly created or updated records against existing table records. When the add-on identifies positive matches during creation or update, you will be prompted to make changes to the record or access merge functionality. See [Automatic Matching](#) for instructions on initiating a matching operation.

After a matching operation, you can view matching related information in the **Consolidated view**. Also, you can access the **Manage group** screen to analyze and take actions on records that are confirmed matches or possible matches (suspects). See [Consolidated view](#) and [Managing groups](#).

Merging

Merging records is the process of combining data values from multiple records into a single, golden record. Records involved in a merge operation that do not become the golden record are called *merged records*. Administrators create *merge policies* that determine how the add-on automatically merges data values. Merge policies can also pre-select data values when you initiate a manual merge.

If a merge policy is configured to happen automatically, the add-on merges records after a matching operation. This happens whether matching was manually or automatically initiated. If matching and merging are set up for automatic execution, these processes will be transparent to you. However, you can view the results in the **Consolidated view**.

When you manually initiate a merge, you use the **Merge view** to select the data values to survive into the golden record. See [Merging Records](#) for details on merging. You can manually initiate a merge:


- From a table's **Actions** menu.
- From the **Manage group** screen.
- When an inline matching operation finds a positive match while creating or updating a record.
(as long as automatic merge isn't enabled)

Matching and Merging Walkthrough

- [Overview](#)
- [Part 1: Initiate matching](#)
- [Part 2: Using the Consolidated view](#)
- [Part 3: Merging suspects](#)

Overview

The following image and accompanying table provide a high-level overview of the matching and merging workflow:



Matching policy initiates

How the add-on is configured determines when the matching policy is initiated. The most common methods are:

- A user manually initiates the matching policy.
- The add-on automatically initiates the matching policy.

See [Part 1 of this walkthrough](#) for more information.

Data is processed


While the add-on is processing data you can choose to run the operation in the

	background and navigate in EBX®, or cancel the operation if needed.
Consolidated view	The Consolidated view displays matching results. You can initiate data stewardship tasks from here. These tasks include: • Validating the results. • Managing groups of candidate records by determining whether to accept suspect records and begin the merge process. See Part 2 of this walkthrough for more details.
Merge process	The merge process allows you to create a golden record by merging the most accurate data from candidate records. See Part 3 of this walkthrough for additional information.


Part 1: Initiate matching

The following steps show how to initiate the matching process manually:

1. Open a table where the add-on is enabled.
2. Determine the matching process scope:
 - To include only specific records in matching, select those records now.
 - Select all records to run matching on the entire table.
3. From the table's **Actions** menu select *Match and Merge > Run match*.



If the matching process takes a while, you can choose to run the process in the background and perform other limited tasks, such as navigate and view your data.



When the matching process completes successfully, the add-on:

- Depending on configuration settings, prompts you on whether to purge matching results.
- Displays the **Consolidated view** or **Manage group** view.

What's next? See [Part 2: Using the Consolidated view](#) for additional details and instructions.

Part 2: Using the Consolidated view

- [Overview](#)
- [Opening the view](#)
- [Matching result stewardship](#)

Overview


The **Consolidated view** is a stewardship view that displays matching-specific metadata. The metadata is linked to each record that was part of a matching process. From this view, you can:

- Get a global view of matching results. Part of the results include: warnings/errors, grouped records, record states, and whether the record was part of the auto-creation process.
- Run a validation report to retrieve information on any groups that contain errors or warnings. The validation report will include any groups flagged for containing suspect records. Suspect records share enough similarity with another record that they are candidates to represent the same entity. You can access and manage the groups of suspects from the validation report.
- Manage matching results by entering the **Manage group** view. Here you can accept or reject candidate records. You can also enter the **Merge view** and reconcile differences between records.

Opening the view

The **Consolidated view** automatically opens after a manually initiated matching completes. However, you can access it at anytime after matching runs on a table.

As shown below, you can open the **Consolidated view** from a table's **Actions** menu by selecting *Match and Merge > Consolidated view*:




Matching result stewardship

Stewardship is required when matching finds suspect records and:

- your merging policy is set to manual.
- results are not definitive enough to trigger an automatic merge.

As a steward in these cases, you will most likely want to verify results, and when required, merge records. The steps below show how to do this using the options in the **Consolidated view**:

1. Find groups of records that require your attention by selecting **Validate** at the top of the view.




2. Analyze matching results:

- a. Locate the **Warnings** description that indicates there are suspect records.

Each entry is a group.

- b. Select a group from the list to enter the **Manage group** view.



3. Manage groups of suspect records:

The idea is to determine whether records in the **Suspects** and **Matches** lists refer to the same entity. To assist in this process the add-on allows you to:

- a. Compare records: Select two or more records in one of the lists and open **Actions > Compare**.

- b. View record metadata to gain insight related to matching such as current state, group, and execution log: Open *Actions > Metadata*.
- c. Once you have made a determination, use the **Actions** menu or select the '-' icon to remove any false positives. Use the **Actions** menu or the check mark icon to confirm suspects and add them to the **Matches** list. When the required suspects are added to the **Matches** list, you can merge the data from these records.

What's next? Use the instructions in [Part 3: Merging suspects](#) to merge records.

Part 3: Merging suspects

- [Overview](#)
- [Merging suspect records](#)


Overview

The **Merge view** allows you to merge data from multiple records into a golden record. This record is the most accurate, up-to-date, and comprehensive representation of an entity. The following section demonstrates how to use the view.

Merging suspect records


To merge data from suspect records and create a golden record:

1. Enter the **Merge view**:
 - a. From the **Manage group** screen (first image) or after inline matching executes (second image), select **Merge**.



b. You can manually select records and from a table's **Actions** menu, select *Match and Merge* > *Merge*.

2. Select the values to persist into the golden record:



1)

Compose the candidate record by selecting the values to persist. Depending on how your administrator has configured the merge feature, the system might pre-select certain data values. Selected values are shaded in gray.

2)

The **Preview** shows what data is included in the golden record. It updates in real-time based on your selections.

3)

Use the **Revert** icon to undo up to 5 of your most recent modifications. Depending on your configuration settings you can use the **Reset** or **Apply merge policy** buttons to reset all of your modifications.

1. Select **Next** to preview results.
2. Select **Merge** to complete the process.

You can select the link to access the record, or **Back to main view** to return to the previous view.

General Tasks and Views

- Automatic Matching
- Evaluating matching
- Searching before taking action
- Manual Matching
- Purging old results
- Managing groups
- Merging Records
- Merging business object records
- Reversing a merge operation
- Consolidated view
- Changing record state
- Viewing add-on metadata

Automatic Matching

If automatic matching is enabled and set to inline, the add-on finds positive results when you create or update a record, the add-on displays the following screen. The list below the screen describes actions you can take here.

Duplicate	Result	Customer ID	Company Name	Website
☒	MATCH	14	Spence Law Offices	http://www.spencelawoffices.com
☐	SUSPECT	50	Anker Law Office	http://www.ankerlawoffice.com

1) Save, Revert, and Delete

You can make changes to record fields and select **Save** or **Save and close**. The add-on

	runs another match. If it finds no duplicates, the record saves normally. The Revert and Delete buttons work as expected and revert any changes or delete the record, respectively. Note that the option to delete is only available when matching is triggered during the process of creating a new record.
2) Merge and Reject all duplicates	You can select the Duplicate checkbox for any positively matched record and select Merge to open the Merge view and merge values. By selecting Reject all duplicates, the new or updated record will be saved in the table as-is. See Merging Records for available actions in the Merge view.
3) Search bar	Use the search bar to locate records in the list of potential matches.
4) Results	Select the details icon in the Result column to display the decision tree and view execution of the logic used to identify the matches. When viewing the decision tree, you can double click a node to view the matching score.

Evaluating matching

- [Overview](#)
- [Running the evaluation](#)
- [Evaluation results](#)

Overview

The add-on's **Evaluate matching** service allows you to check the outcome of a matching operation against expected results. This can help you fine tune configuration settings to meet your specific needs.

This service provides the following benefits:

- Quick results
- Ability to test different policies without exiting the service
- Uncover points of matching execution failure
- Upon completion, view details that show how the algorithms score and a match/no match was reached
- With administrative credentials, directly access configuration settings without leaving the screen

Running the evaluation

To run the service:

1. Select two records from a table with a configured matching policy.
2. From the table's **Actions** menu, select **Evaluate matching**.
3. Use the **Matching policy** menu to select the policy you want to test.

With sufficient permission, you can access the matching configuration and make changes.

4. Select **Launch** to run the evaluation.

After running, the results display at the bottom of the screen.

Evaluation results

The following image and list describe results of the matching process:

Note:

For more details on the matching process, see [Understanding matching operation processing](#).


Please select the matching policy to evaluate the matching result of the two selected records.

Matching policy  **2**

1 - 2 of 2    

ID	First Name	Last Name	DOB	Street	City	Country	Tower	Number	C
2	John	Smithe	12/11/1980	1234 Harbor Dr	Athens	Greece	1	1000	N
3	John	Smith	12/11/1980	1234 Harbor St	Athns	Greece	1	1000	N

Matching result



```

graph LR
    1((1)) --- 2((2)) --- 3((3)) --- 4((4))
    1 --- Start[Start]
    2 --- Pre-processing[Pre-processing  
The process to regroup similar records before applying matching policy to identify duplicates.]
    3 --- Decision-tree-execution[Decision tree execution  
Determines which data values are compared and how to compare them in a matching operation.]
    4 --- Decision-tree-result[Decision tree result  
3 MATCH  
See details »]
  
```

1 **SIMILAR**
[See details »](#)

3 **MATCH**
[See details »](#)

[Launch](#) [Close](#)

1) Pre-processing

This step groups records that can be considered as similar and returns one of the following values: Similar, Not Similar, and Inconsistent. Click **See details** to see the tables, fields, and strategies used in the evaluation. Some matching policy configurations might cause inconsistent results in the pre-processing phase of matching. When records are considered as being Not Similar, the process stops at step 2, the records do not pass the pre-processing step, and are prevented from evaluation using the matching logic in the decision tree. If you notice records that should have made it past pre-processing but have not, you can:

	- Analyze the fields configured to use matching. Depending on your data, too many irrelevant, or too few important configured fields can cause the system to view records as not similar. - Ensure that the correct weight is used for fields used in matching. Assign higher weights to the more important fields. - Verify that you are using the recommended search strategies for each field's data type.
2) Matching policy	Change policies or configuration settings on the fly and re-test. Select the  icon to access policy settings.
3) Final result	The final matching result (Match/No Match/Suspect). Select details to view a map of the decision tree execution. Double-click a data comparison node to view details about how the records were processed.

Searching before taking action

- [Overview](#)
- [Using search before create](#)

Overview

The **Search before create** service can help prevent the duplication of existing data. It allows you to determine whether the values from a record you want to create potentially match existing records. When searching, you select the matching policy used to compare data values. The selected policy determines which fields are available for the search, and the criteria used to determine potential matches.

Note:

You cannot use fields from associated tables with this service when matching is configured to use relationships.

Using search before create

Search for potential duplicates prior to creation using the following steps:

1. From a table's **Actions** menu, select *Match and Merge > Search before create*.
2. Select the desired matching policy to determine which fields are available to search and how the add-on determines potential matches.
3. Fill the available fields with information from the record you want to create and select **Search**.
4. After the search completes you have the options to select:
 - **Reset** to clear the search form's fields.
 - **Update record** icon in the **Potential matches** panel to change values in an existing record.
 - **Create a new record** to add a new record to the table. In the creation form that displays, any fields that included searched values will be populated with these values.
 - The details icon (in the **Results** column) to view the decision tree and see an overview of the logic used to determine positive matches. When viewing the decision tree, you can double click a node to view the matching score.

Search before create

Customer Contacts table

Choose the desired matching policy and enter values for the new record to check for duplicates.

Matching policy:

First Name:

Last Name:

Potential matches

1 potential match have been found for this record.

Action	Result	Customer Contact ID	Company	First Name	Last Name	Primary phone number	Secondary phone number	Group
	MATCH	1	100	Fred	Myers	555-555-1212	555-656-2223	

Manual Matching

When you manually run a matching operation, you must select one or more records. The operation compares the selected records with each other and with all the other records in the table. The add-on creates groups of records that match, are suspect matches, or are golden. Each group is assigned a number, contains at least one selected record, and can include unselected records. Note that unselected records are *only* compared with selected records. Unselected records might match each other, but they are only included in a group when they match a selected record.

Consider the following image and description that shows the result of a matching operation where only the first 3 records were selected.

1	pk	color	Group	State
☒	1	red	1	
☒	2	blue	2	Golden
☒	3	green	3	
☐	4	red	1	
☐	5	green	3	
2	6	yellow		
☐	7	yellow		

1)

The add-on created the following groups:

- Group 1 includes records 1 and 4 (values of red).
- Group 2 includes only record 2 with a value of blue as it has no matches/suspects and is considered golden.
- Group 3 includes records 3 and 5 (values of green).

2)


Even though records 6 and 7 have matching values of yellow, the add-on did not group them since they were unselected and not compared against each other.

Attention:

Performance and result accuracy can be improved by purging previous results before manually executing a matching operation. Use the **Purge old results** service to accomplish this.

To run a manual match on a table that has an active matching policy:

1. Set the matching scope by selecting one or more records.
2. Open the **Actions** menu and select *Match and Merge > Run match*.



3. In the **Run match** screen, select the matching policy to use from the list of active policies and run the match.

If you executed the match:

- On an entire table or a selection of multiple records, the **Consolidated view** displays. See [Consolidated view](#) for more information about actions available.
- With only one record selected, the **Manage group** screen displays. See [Managing groups](#) for more information about actions in this screen.

Note:

While the matching process is running, the add-on displays its progress. At any time you can cancel the process, or navigate in EBX® to view other data. When the matching process completes, you can view the results in the **Consolidated view**.

Purging old results

The **Purge old results** service removes outdated matching metadata generated by the add-on. This helps improve matching performance and reduces storage usage.

Note that matching metadata is stored in a technical dataset created by the add-on. A separate dataset is created for each table where matching is executed. The following naming convention is used for datasets: `<datasetName>\_<tableName>\_MDS`.

When you run the **Purge old results** service:

- Redundant suspect relationships are removed between records that were flagged as false-positive. This impacts the **MatchingResult** and **FalsePositiveDecision** metadata tables.
- Outdated matching results where the source and target record appear in multiple matching processes are cleaned. This impacts the **MatchingSelection** and **MatchingResult** metadata tables.
- The **MatchingSelection** table is truncated.
- Records that are orphaned for any of the following reasons are removed from the **RecordMetadata** metadata table:
 - If record was, but is no longer linked to a golden record (for example if the golden was deleted).
 - If a record was deleted in the main table, but still exists in the **RecordMetadata** table.

Note:

If this service is run on a Business Object, all related child metadata tables are also purged.

Managing groups

During matching, the add-on groups records that meet the criteria to be considered positive matches. The add-on also identifies records that are suspected of matching records in a group. When you manage a group of records, you can compare records, view record metadata, accept

suspect records as duplicates, and reject false positive results. You can also enter the **Merge view** to merge data into a golden record.

If enabled by an administrator, an additional **Match type** column is displayed. This column identifies records that are included in the group because they are a transitive match. Transitive matching refers to the inclusion of records based on indirect relationships. For example, if Record A matches Record B, and Record B matches Record C, then Record A and C are grouped together due to their transitive connection through B. The column identifies records with the following match types:

- **Direct:** The record was matched directly to the group's survivor.
- **Transitive:** The record was matched indirectly through another matched record.
- **N/A:** The group's survivor record; match type does not apply.

After a matching operation, you can open the **Manage group** view:

- By selecting a record and then accessing the **Manage group** service from a table's **Actions** menu.
- From the **Consolidated view** or **Metadata view** by selecting the **Manage group** icon next to a group.

The following image and table describe available actions:

1) Actions menu, View menu, and filter

2) The **Actions** menu allows you to:

State	Customer Contact ID	Company	First Name	Last Name
Unset	1	540	Jon	Smith
Unset	2	325	John	Smith
Unset	3	454	John	Smithe

3) Close

4) Merge

- Select and remove records from the list of matches by selecting *Matches > Remove*. Any removed records are categorized as false positive matches and are not grouped with a record's suspects or matches. See the bullet point below on metadata for additional information about false positive categorization.
- Accept or reject suspect records as matches using the corresponding option under **Suspects**. When you confirm that a record is a positive match, the add-on moves it to the list of matches where it can be used in a merge operation. If you reject a suspect record, the add-on marks it as a false positive and will no longer include it in groups with the current records.
- Compare two selected records using the **Compare** service.
- Use the **Evaluate matching** service to see how two records were evaluated in the matching process.
- Use the **Unmerge** option to revert a merge operation.
- Use the **Metadata** service to view a selected record's matching-specific metadata. When viewing metadata, you can use the **Identified as false positives** tab to see records that were manually identified as not matching with the selected record. False positive records are no longer grouped with a record's suspects or matches. You can

use the **Actions** menu when viewing metadata to remove the false positive categorization to reverse this behavior. Note that in some instances, merged records marked as false positive can be re-included in groups. In these cases the add-on uses the



icon as a designation. You can hover your mouse over the icon to view list of related false positives.

View menu: By default, any recommended or favorite view is applied. Otherwise, the table's default view is used. The **View** menu allows you to also select from any predefined table views. This allows you to create specific tabular views optimized for performing matching-related tasks. Note that the menu displayed in the add-on does not provide the option to create a new view. Instead, create these views from the standard EBX® tabular view by opening a table's **View** menu selecting **Create a new view**.

Click the **Filter** icon to activate filters that: hide any merged records in the view, or to hide all fields except those configured to participate in matching.

- Hide any merged records in the view.
- Display only fields configured to participate in matching.
- Display only records that are transitive matches. Note that this option is only visible when the matching policy

	configuration has the Show match type property enabled. The Display records menu allows to switch between vertical and horizontal display of the records. The view mode you select here is also maintained when navigating to the Merge screen.
2) List of matches	The list of records identified as matching. Records are considered matches if they meet the criteria set out in the matching configuration, or if you confirm a record in the Suspects list. Use the Reject icon to remove a record from the list. Double-click a record value to open the record in the details view.
3) List of suspects	This list contains records that were determined to be suspects. A record is considered a suspect when it meets the criteria defined in a matching policy. The icons displayed at the top of the list and across from each record allow you to: reject all records, reject individual records, or accept individual records. Double-click a record value to open the record in the details view.
4) Merge button and Set as golden icon (not shown)	The Merge button opens the Merge view and makes all of the records in the Matches pane available for a manual merge. The  displays when the Matches group contains only one record and its state is not yet golden. Selecting this icon sets the record to golden. If this button is enabled and there

are suspects, setting the record to golden will automatically reject the others as suspects.

Attention:


The **Manage group** view is slightly different when matching is performed on business objects. The **Actions** menu combines options for the **Matches** and **Suspects** pane. Additionally, you use the **Related tables** menu to toggle display of the business object's related tables.

Merging Records


The process of merging data allows you to select the fields in records that contain the most accurate data and *merge* them into a golden record. A golden record represents the most accurate and up-to-date version of data.

To merge records:

1. Access the merge view in one of the following ways:
 - When in the EBX® table view, select two or more records to merge:



- When using the add-on's **Manage group** service:



- When inline matching is active and finds potential matches for a newly created or updated record.

The screenshot shows a web interface for managing customer records. At the top, there's a 'Main' tab. Below it, the 'Customer ID' is 876. The 'Company Name' is 'Bobs Books' and the 'Website' is 'www.bobsbooks.com'. There are three buttons: 'Save' (blue), 'Revert' (light blue), and 'Delete' (light blue). Below these buttons, a message states '1 potential match have been found for this record.' Underneath this message are two buttons: 'Merge' and 'Reject all duplicates'. At the bottom, there is a table with columns: 'Duplicate', 'Result', 'Customer ID', 'Company Name', and 'Website'. The table contains one row with a checkbox in the 'Duplicate' column, the word 'SUSPECT' in the 'Result' column, and the values '857', 'Bob s books', and 'www.booksbybob.com' in the other columns.

Duplicate	Result	Customer ID	Company Name	Website
☐	SUSPECT	857	Bob s books	www.booksbybob.com

- Click to select the fields you want to merge into the golden record. If not already selected automatically, selecting a record's primary key selects the entire record. You may want to select the most accurate record first to act as a baseline. Then you can fine tune values by selecting other fields. The choices you make are reflected by the preview at the bottom of the page. Once satisfied, click **Next**. Note that if a merge policy is configured, some values will be pre-selected. You can still change these values.

Merge

1. Customers Step 1/3 Reset

Please select the values that you want to consolidate in the golden record from duplicate records below.

View [Default view] Display records Horizontally

Customer ID	Company Name	Website
1	My Book Oasis	www.mybookoasis.co...
2	The Book Oasis	www.bookoasis.com

Preview

Customer ID	Company Name	Website
1	The Book Oasis	www.bookoasis.com

Cancel Next

1) Undo and Reset or Apply merge policy

You can undo your last five actions, or reset all selected values to their original selections when entering the **Merge view**.

2) Record display option

Use **View** menu to select a table view defined in EBX®. Use **Display records** to set your preferred record display (vertical or horizontal).

3) Data selection area

Select the data values that should be survived into the golden record.

4) Preview record

The **Preview** shows the merge result with the currently selected values.

1. This step depends on your data structure and configuration settings:
 - If this merge will not impact other tables, or your configuration is set to ignore relationship dependencies, the add-on presents you with the **Preview** page. See the next step for more information.
 - If other tables can be impacted by this merge, you are presented with a page to select the relationships to update. The add-on pre-selects values based on the selected golden record when entering the **Merge view**. If you do not have sufficient permissions to display the dependency, the add-on automatically makes the selection based on the record considered the golden record.
2. When you reach the **Preview** step, the add-on provides you an overview of what the merged record will look like. Select **Merge** to complete the process, or use the navigation dropdown to access previous steps to make changes.

Note:

You cannot update records that are in the **Merged** or **Deleted** states.

Merging business object records

- [Overview](#)
- [Merging business object records](#)
- [Business object merge UI overview](#)
 - [Main table merge screen](#)
 - [Related and child table merge screen](#)
 - [Merge preview screen](#)
 - [Merge summary screen](#)

Overview

The process of merging data allows you to select the fields in records that contain the most accurate data and *merge* them into a golden record. A golden record represents the most accurate and up-to-date version of data. You can perform merge tasks for the multiple tables that comprise a business object.


Note:

When performing a manual merge on business objects, you are limited to 3 levels of tables. In other words, if the main table in the business object is considered the parent table, you can manually merge records from related tables to the grandchild level.

Merging business object records

To merge records:

1. Access the merge view in one of the following ways:
 - When viewing the business object's main table in the EBX® table view, select two or more records to merge:



- When using the add-on's **Manage group** service, the **Merge** button is located in the bottom-right corner of the screen.

2. Choose values for the main table:

Click to select the fields you want to merge into the golden record. The choices you make are reflected by the preview at the bottom of the page. Once satisfied, click **Next**. Note that if a merge policy is configured, values will be pre-selected. You can still change these values manually.

See [Business object merge UI overview](#) for an overview of the available options on this screen.

3. Choose values for related tables:

Note:

See [Business object merge UI overview](#) for an overview of the available options on this screen.

- a. Click **+ New golden** located next to the **Result preview**. You cannot select values to merge until completing this step.

Note:

You can create multiple golden records. However, you can only select values from the same record once. After all of the available values have been used, you can no longer create another golden.

- b. Add values to the golden record by:

- Ticking the box to enable value selection for the records you want to merge. The values from the first record you select are automatically added to the new golden.
- Choosing values from the selected records.
- Optionally, repeat the above steps to create another golden record.

- c. If required, merge records from child tables:

- In the result preview, select the



icon.

- Create a new golden record and add the values you want to merge into it.
- If there are additional child tables to merge, use the button at the bottom of the screen to navigate between them.

- d. After completing the above steps for all related tables, click **Next** to see the merge preview. If you chose not to merge some of the records, a warning pop-up displays the table and number of its records that will no be merged.

4. When you reach the **Preview** step, the add-on provides you an overview of what the merged record will look like. Select **Merge** to complete the process, or navigate to previous steps to make changes.

After the operations complete, the **Summary** screen informs you of whether the merge was successful.

Note:

You cannot update records that are in the **Merged** or **Deleted** states.

Business object merge UI overview

The following sections highlight some of the UI features when merging a business object:

- Main table merge screen
- Related and child table merge screen
- Merge preview screen

Main table merge screen

The following image and table highlight some of the options and functionality when merging the business object's main table:

Merge view > Merge Person

View ▾

Selected records

ID	First Name	Last Name	DOB	Street	City	Country
2	John	Smithe	12/11/19...	1234 Harbor Dr	Athens	Greace
3	John	Smith	12/11/1980	1234 Harbor St	Athns	Greece
4	Jon	Smith	12/11/1980	1234 Harber ...	Athens	Grece

Result preview

ID	First Name	Last Name	DOB	Street	City	Country
[auto generated]	John	Smith	12/11/1980	1234 Harber Dr	Athens	Greece

Buttons: Reset, Undo, Cancel, Next

1) Reset, Undo, and Apply merge policy	Select Reset to return all selected records to their original values. Select Undo to revert the most recent value selection. The undone change is highlighted in the Result preview. Apply merge policy is only available if a default merge policy is selected and automatically selects values for the golden record according to the merge policy.
2) Data selection area	Select the data values to survive into the golden record.
3) Merge progress tracker	This component allows you to track your progress during the merge process. It displays the current step, including sub-steps if applicable. Hover your mouse over a step to display a tooltip.
4) Result preview	The golden record result based on the currently selected values to merge.

Related and child table merge screen

The screens for merging relate tables and child tables are mostly similar with a few minor differences. The following image and table highlight some of the merging options and features specific to related tables and child tables:

Note:

To enable selection of values to merge, begin by selecting **New golden**.

Merge view > Merge related objects: Company

View ▼ Reset Undo

Company

☒ 1	Person	ID	Company Name	Street	City	Country	Tower	Number	Child records
☒ 2	2	1	X Large Group						4
☐ 1	3	2	Murphy N Sons						3

Result preview + New golden 3

☐ 2	Person	ID	Company Name	Street	City	Country	Tower	Number	Child records
☐ 1	[Link...]	[a...]	Murphy N Sons						0 (0/3)
☐ 2	[Link...]	[a...]	X Large Group						0 (0/4)

5 Back Cancel Next

1) Record selection

Selecting a record's checkbox allows you add its values to a new golden record. Once values have been added to a golden record, a number displays to indicate which record in the **Result preview** section contains its values.

2) Result preview

Use the icon to remove a golden record. This will free up the values to be used in other goldens.

3) + New golden

Click this button to create a new golden record. Note that you cannot select values from the records above until a golden record is available to accept them.

4) Child records (related table level only)

This displays at the business object's related table level and shows the number of child

	records available to merge, and that have been merged. Use the  icon to access merge options for records in the child level.
5) Navigation buttons	Use the Next and Back buttons to navigate between the business object's levels. When merging at the related table level, these buttons navigate between related tables. If you reach the final related table, and have skipped merging any child level records, a pop-up displays to ensure you want to continue. To access the child table level use the icon in the Child records column (described above). You can then navigate through and merge the child tables. When you reach the last child table, the Next button takes you to the Preview screen.

Merge preview screen

The following image and table highlight some of the options specific to previewing the merge results:

Merge view > Preview

1

2

2.1

2.2

2.2.1

2.2.2

3

4

1

2

3

ID	[auto generated]
First Name	John
Last Name	Smith
DOB	12/11/1980
Address	1234 Harbor St Athens Greece
Address Details	ID Company Name Street City Country Tower Number [auto generated] Murphy N Sons Customer ID Customer Name [auto generated] A2 Banking [auto generated] X Large Group
Gender	M
Status	Employeeed
History	
Age	43
Source	ERP
Company	

Back Cancel Merge

1) Main record values	Values from the business object's main table display first.
2) Related record values	Records from related tables display in the top level tabular view.
3) Related record children	Child records of related tables are nested here.

Merge summary screen

This screen provides a summary of the merge operation that was just completed and allows you to access newly created golden records.

Reversing a merge operation


You can rollback a merge operation using the **Unmerge** service. This service reverts the last merge operation for the selected golden or merged records. Records will remain duplicate.

To reverse a merge operation:

Note:

You can perform the same action described below from the **Manage group** screen.


1. Select either a golden or merged record.
2. From the table's **Actions** menu, select *Match and Merge > Unmerge*.



Consolidated view

The **Consolidated view** allows you to view add-on related metadata applied to a table. Note that records only display here if they have matching related metadata.

The following image and list describe features and available actions:



1) Group and Manage group link

The add-on groups records positively identified as matching records. This field identifies which group each record has been placed in.

Select the icon next to a group to open the **Manage groups** screen. This screen allows you to compare records and confirm or reject records as duplicates.

2) State

When a record has taken part in a match or merge operation, the add-on assigns it state. The state identifies the result of these operations.

3) Auto created

Indicates whether the corresponding record was automatically created during match or merge operation.


Changing record state

When the add-on completes match and merge operations, it sets a state for each evaluated record. Possible states include:

- **Golden:** The singular record that most completely defines a business entity.
- **Merged:** A record that has been included in a merge operation, but does not meet the requirements for a golden record.
- **Unset:** A record that has null matching state.
- **Pending approval:** after an unmerge operation, all records that were part of the initial merge operation are set to pending approval.
- **Deleted:** Auto-created golden records that are no longer targeted by other records, or were merged with other records are deleted.

If for some reason you want to change a record's state, you can use the **Change state** service. To use this service:

1. Select one or more records to change their state.
2. From the table's **Actions** menu, select *Match and Merge > Change state*.



3. Choose the target state and select **Launch service** at the bottom of the screen.

Change state

?

Please select the target state value of your selected records from values below.

Target state

<unset>

Golden

Cancel

Launch service

Note:

You cannot change the state for records that are in the **Deleted** state.

Viewing add-on metadata


- [Opening the view](#)
- [Viewing golden record lineage](#)
 - [Superseding survivorship](#)

Opening the view


The add-on uses metadata to track records where the add-on is active. The metadata includes information about record state, groups, and target records. Additionally, the metadata tracks matching execution. To view record metadata:

1. Select a record and from the **Actions** menu, select *Match and Merge* > *Metadata*. As shown in the images below you can access the view from several locations:


Table view:




Record details view:



Stewardship screens such as **Manage group**:



2. From the metadata screen you can perform the actions described below:



1) Execution log, Lineage, and Identified as false positives tabs

These tabs contain the following information:

- **Execution log:** Displays the most recent matching execution date and time. Additionally, it shows the most recent stewardship decision, date and time, and user that performed the action.
- **Lineage:** Shows any lineage related to a golden record including values from

	records that were merged into this record. • Identified as false positives: Lists any records that were rejected as suspects or matches with this record. The false positive records will no longer be grouped with this record in future matching operations.
2) Record icon	Use this icon to open the record in edit mode.
3) Group icon	Select this icon to open the Manage group screen. You can confirm suspect records from this screen.

Viewing golden record lineage

The **Metadata** screen contains the **Lineage** tab which allows you to view the values from records that were merged into a golden record. To ensure the most accurate lineage is tracked, you must enable EBX® history in the Data Model Assistant.

Access the **Lineage** tab in the following ways:

- Select a record and from a table's **Actions** menu choose **View lineage**.
- When in a record's details view, open the **Actions** menu and choose **View lineage**.
- When viewing a record's metadata, select the **lineage** tab.

The screenshot displays the 'Lineage' tab in the TIBCO EBX interface. It shows two main sections: 'Merged records' and 'Golden record'. The 'Merged records' section contains a table with columns: id, State, First name, Last name, Primary Phone Number, and Secondary Phone Number. The 'Golden record' section contains a similar table. Red circles 1-4 highlight specific features: 1) The 'Merged records' table, 2) Highlighted values in the 'Golden record' table, 3) The 'Actions' menu options, and 4) The 'Golden record' table.

id	State	First name	Last name	Primary Phone Number	Secondary Phone Number
344	Merged	An	Fritz	609-228-5266	609-228-5265
336	Merged	An	Fritz	555-555-1212	

id	State	First name	Last name	Primary Phone Number	Secondary Phone Number
342	Golden	An	Fritz	609-228-5266	609-228-5265

1) Merged records

The list of merged records includes any records that were part of a merge operation to create the golden record. This is true even if no values from a given record were merged into the golden record.

2) Highlighted values

These values were merged into the golden record. Hovering your mouse over a highlighted value shows the merge function and condition used to determine that this value should be merged into the golden.

3) Actions menu options

The **Actions** menu in the **Lineage** tab includes the following options:

- **Supersede survivorship:** This allows you to override merge decisions to change, without unmerging, values in the golden record. For more information, see [Superseding survivorship](#).
- **Unmerge:** This service is like performing an undo action on the most recent merge

	operation that created the selected golden record.
4) Golden record and highlighted values	This group shows the golden record that resulted from a merge. Highlighted values are those that survived from the original record. In the example shown, the golden record kept its original values from the id, First name, and Last name fields.


Superseding survivorship

In some instances incorrect values are merged into golden records. When this happens, you can run the **Supersede survivorship** service to change the values in the golden record. This service allows you to reverse merge decisions on specific attributes without using the **Unmerge** service to completely revert the merge process.


Depending on how an administrator has configured trusted source options, the decision you make to override a previous merge can impact future automatic merges, or merge recommendations during manual merge. For example, you might override a value so that it comes from an HR system over an IT system. The next time the record is involved in a merge operation, the add-on can automatically prioritize the value coming from the HR system. If you are in an administrative role and would like to learn more about using and configuring trusted sources, see [Setting trusted sources](#).

To supersede survivorship:

1. Open the **Metadata** screen for a golden record using one of the following methods:
 - a. From a table view, select a golden record and from the **Actions** menu, select *Match and Merge > Metadata*.
 - b. From the **Manage group** view, select a golden record and from the **Actions** menu, select **Metadata**.
 - c. Alternatively, open a golden record from the **Manage group** view and select **Metadata** from the **Actions** menu.



2. Run the **Supersede survivorship** service:
 - a. Select the **Lineage** tab.
 - b. From the **Actions** menu, select **Supersede survivorship**.



3. Select the values to survive into the golden record. Selected values are shaded in gray.
4. Select **Next**.
5. Double-check the results and select **Supersede** to perform the operation and view the results page.

Release Notes

- [Version 6.2.3](#)
- [All Release Notes](#)

Version 6.2.3

- [New features](#)
- [Changes in Functionality](#)
- [Upgrade notes](#)
- [Changes to third-party libraries](#)
- [Closed Issues](#)
- [Known Issues](#)

New features

This release contains the following enhancements:

Viewing record lineage:	• Record lineage is now available in toolbars and workflows. This provides the convenience of single-click access to golden record lineage via a button in a custom record form. • The View lineage option is now accessible directly from the Actions menu when viewing record details. This prevents you from having to navigate to the Metadata view to access lineage.
Performance enhancements	Updates in this release have reduced redundant matches when processing tables that have many duplicate records. This speeds up matching, grouping, and merging,

	and results in fewer records in the Matching results table.
Automated purging	The add-on now supports automated purging of outdated matching metadata through the EBX® Task scheduler. This enhancement reduces manual maintenance and helps prevent performance degradation over time. For more information, see Scheduling matching metadata purges.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Note:

As the add-on UI is under development, some images in the documentation might not exactly match those in the UI.

Changes in Functionality

This release contains no changes in functionality.

Upgrade notes

Take the following into account if you upgrade TIBCO EBX® Match and Merge Add-on:

- When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.
- If upgrading from any version prior to 6.2.0, a performance improvement was made to better balance the weights of fields that participate in matching. The automatic balancing occurs

during the pre-processing phase. This helps mitigate issues that might occur when fields include many values. As a result of this update, matching operation results might be impacted. To mitigate any negative impacts, you can test and adapt your existing weights to achieve the desired matching outcome. Alternatively, you can set the `ebx.clustering.balance` property to `false` in the `ebx.properties` file to restore the legacy behavior.

Changes to third-party libraries

The following third-party libraries were updated:

- Jackson Core was updated to version 2.21.1.
- Spring Framework was updated to version 6.2.14.
- Jakarta Servlet was updated to version 6.0.

Closed Issues

This release contains the following closed issues:

- **MAME-9523** When matching with business objects, group management is inaccessible after deleting the center record.
- **MAME-9904** Records matched through transitive relationships are not grouped correctly when matching is executed in batches.
- **MAME-9914** The tooltip description for the **Any duplicates** merge mode should be rephrased and clarified.
- **MAME-9930** An error occurs when importing the add-on's metadata (_MDS dataset).
- **MAME-9933** The add-on's logs are throwing an error in EBX® Container Edition.
- **MAME-9952** When a matching policy is set to execute after submission, it is not triggered when creating a new record from a workflow.
- **MAME-9959** When the merge policy is set to **Auto create new golden** and the mode is **Only duplicates**, any customized source value that should only apply to new golden records is getting incorrectly assigned to unique golden records.

Known Issues

This release contains the following known issues:

- The following workaround is for a known issue in EBX® (CP-27860). When child dataspace or snapshots are created before enabling the EBX Match and Merge Add-on on a data model, some operations can fail after the add-on is enabled. For instance, child dataspace merge might not work correctly, or snapshots might have errors.

Workaround: Disable the add-on before performing the dataspace merge, or creating a snapshot; enable the add-on after completing the operation:

- a. Open the data model, navigate to *Configuration > Add-ons*.
 - b. To disable the add-on, tick the box next to **Match and Merge** and select **Delete** from the **Actions** menu.
 - c. Publish the data model.
 - d. After performing the desired operations, enable the add-on by re-creating the **Match and Merge** record in the **Add-ons** table.
 - e. Publish the data model.
- The **Evaluate matching** service returns inconsistent results for related matching fields that have the DateTime data type.
 - It is not possible to accept or reject suspects in groups that contain a large number of records.
 - Foreign key alignment of recursive foreign keys is not possible.
 - Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
 - A manual merge operation cannot be performed on more than 400 records.
 - When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
 - The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
 - In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
 - Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.

- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

All Release Notes

- Version 6.2.3
 - New features
 - Changes in Functionality
 - Upgrade notes
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 6.2.2
 - New features
 - Changes in Functionality
 - Upgrade notes
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 6.2.1
 - New features
 - Changes in Functionality
 - Upgrade notes
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 6.2.0
 - New features
 - Changes in Functionality

- Upgrade notes
- Changes to third-party libraries
- Closed Issues
- Known Issues
- Version 6.1.3
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 6.1.2
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 6.1.1
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 6.1.0
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues

- Version 3.6.1
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.6.0
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.5.0
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.4.1
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.4.0
 - New features
 - Changes in Functionality
 - Changes to third-party libraries

- Closed Issues
- Known Issues
- Version 3.3.5
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.3.4
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.3.3
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.3.2
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.3.1

- New features
- Changes in Functionality
- Changes to third-party libraries
- Closed Issues
- Known Issues
- Version 3.3.0
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.2.1
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.2.0
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.1.1
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues

- Known Issues
- Version 3.1.0
 - New features
 - Changes in Functionality
 - Changes to third-party libraries
 - Closed Issues
 - Known Issues
- Version 3.0.1
 - Bug Fixes
- Version 3.0.0
 - Release overview

Version 6.2.3

New features

This release contains the following enhancements:

Viewing record lineage:	• Record lineage is now available in toolbars and workflows. This provides the convenience of single-click access to golden record lineage via a button in a custom record form. • The View lineage option is now accessible directly from the Actions menu when viewing record details. This prevents you from having to navigate to the Metadata view to access lineage.
Performance enhancements	Updates in this release have reduced redundant matches when processing tables that have many duplicate records. This

	speeds up matching, grouping, and merging, and results in fewer records in the Matching results table.
Automated purging	The add-on now supports automated purging of outdated matching metadata through the EBX® Task scheduler. This enhancement reduces manual maintenance and helps prevent performance degradation over time. For more information, see Scheduling matching metadata purges.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Note:

As the add-on UI is under development, some images in the documentation might not exactly match those in the UI.

Changes in Functionality

This release contains no changes in functionality.

Upgrade notes

Take the following into account if you upgrade TIBCO EBX® Match and Merge Add-on:

- When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

- If upgrading from any version prior to 6.2.0, a performance improvement was made to better balance the weights of fields that participate in matching. The automatic balancing occurs during the pre-processing phase. This helps mitigate issues that might occur when fields include many values. As a result of this update, matching operation results might be impacted. To mitigate any negative impacts, you can test and adapt your existing weights to achieve the desired matching outcome. Alternatively, you can set the `ebx.clustering.balance` property to `false` in the `ebx.properties` file to restore the legacy behavior.

Changes to third-party libraries

The following third-party libraries were updated:

- Jackson Core was updated to version 2.21.1.
- Spring Framework was updated to version 6.2.14.
- Jakarta Servlet was updated to version 6.0.

Closed Issues

This release contains the following closed issues:

- **MAME-9523** When matching with business objects, group management is inaccessible after deleting the center record.
- **MAME-9904** Records matched through transitive relationships are not grouped correctly when matching is executed in batches.
- **MAME-9914** The tooltip description for the **Any duplicates** merge mode should be rephrased and clarified.
- **MAME-9930** An error occurs when importing the add-on's metadata (_MDS dataset).
- **MAME-9933** The add-on's logs are throwing an error in EBX® Container Edition.
- **MAME-9952** When a matching policy is set to execute after submission, it is not triggered when creating a new record from a workflow.
- **MAME-9959** When the merge policy is set to **Auto create new golden** and the mode is **Only duplicates**, any customized source value that should only apply to new golden records is getting incorrectly assigned to unique golden records.

Known Issues

This release contains the following known issues:

- The following workaround is for a known issue in EBX® (CP-27860). When child dataspace or snapshots are created before enabling the EBX Match and Merge Add-on on a data model, some operations can fail after the add-on is enabled. For instance, child dataspace merge might not work correctly, or snapshots might have errors.

Workaround: Disable the add-on before performing the dataspace merge, or creating a snapshot; enable the add-on after completing the operation:

- Open the data model, navigate to *Configuration > Add-ons*.
 - To disable the add-on, tick the box next to **Match and Merge** and select **Delete** from the **Actions** menu.
 - Publish the data model.
 - After performing the desired operations, enable the add-on by re-creating the **Match and Merge** record in the **Add-ons** table.
 - Publish the data model.
- The **Evaluate matching** service returns inconsistent results for related matching fields that have the DateTime data type.
 - It is not possible to accept or reject suspects in groups that contain a large number of records.
 - Foreign key alignment of recursive foreign keys is not possible.
 - Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
 - A manual merge operation cannot be performed on more than 400 records.
 - When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
 - The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
 - In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
 - Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.

- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 6.2.2

New features

This release contains the following enhancements:

- Null value management:
 - Previously, a field's null value management settings were applied anytime the field was used in a decision tree's comparison node. To provide more flexibility, comparison nodes now include settings for each field to determine the outcome when one or both of the compared values is null. This makes it possible to apply different matching behaviors for the same field in different parts of the decision tree. For example, one node can require an exact match, where both records must contain a value, while another node can use a "not mismatched" strategy, which considers records matching unless their values are explicitly different. Note that these new settings apply only to the decision tree processing phase of matching. When upgrading from a previous version, existing field configurations default to **No match** for the null value options.
 - A new **Do nothing** option was added to the **Matching fields** tab's options. This option ensures that null values are not considered as similarity criteria during the pre-processing clustering phase. When null values are ignored, matching performance and quality can improve as it can prevent unnecessary decision tree computations. For instance, it can prevent false positives from being passed to decision tree processing for evaluation.
- The following new options were added for specifying when golden records are automatically created for business objects:
 - **Only duplicates** (available for a business object parent entity only): A new golden record is automatically created only when the add-on identifies a positive match between records. With this option selected, child entities must use the **Any duplicates** option described below.
 - **Any duplicates** (available for child entities only): If the parent has positive matches, a new golden is created for the parent and its children. If matches are found for children of the same parent, but not the parent, a new golden is created for the child only.

- This release includes the option to show whether records were matched through a transitive match. Transitive matching refers to the inclusion of records based on indirect relationships. For example, if Record A matches Record B, and Record B matches Record C, then Record A and C are grouped together due to their transitive connection through B. Administrators can configure display of this information using the new **Show match type** property, located on each matching policy. When enabled, users will have a new **Match type** column in the **Manage group** screen that shows whether the match type is **Direct**, **Transitive**, or **N/A** for the survivor record.
- Performance improvements: This release includes multiple backend updates and enhancements to improve the add-on's performance.
- Improved exception handling:
 - Workflows that include matching operations are now stopped when an exception occurs in the add-on. An error message detailing the issue displays in the workflow UI.
 - When performing a grouping operation, an exception is thrown when the center record cannot be found. Previously, only a `NullPointerException` was thrown. Now, an error message is logged and shown in the UI. The message includes the group Id to support location and correction of the issue.
- Logging enhancements:
 - Warnings for `SurvivorServiceImpl#getFieldSurvivorsByMergeFunction` could be repeatedly logged, leading to a performance cost and difficulty in parsing the logs. To alleviate this, the warnings are consolidated and logged after longer increments.
 - A log entry is now created when pausing and resuming the matching trigger. The log includes information on the affected table.
 - Logs created when merged data includes many null fields were condensed to improve readability.
- API enhancements: The simulate match API previously returned a list of `MatchingEngineResult` objects that contained primary keys for the target and pivot records. This required additional resources to look up the record's corresponding `Adaptation`. Now the `Adaptation` is returned instead of the primary key.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Note:

As the add-on UI is under development, some images in the documentation might not exactly match those in the UI.

Changes in Functionality

In addition to existing functionality, the **Purge old results** service removes any records that are orphaned for any of the following reasons:

- If record was, but is no longer linked to a golden record (for example if the golden was deleted).
- If a record was deleted in the main table, but still exists in the **RecordMetadata** table.

Upgrade notes

Take the following into account if you upgrade TIBCO EBX® Match and Merge Add-on:

- When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.
- If upgrading from any version prior to 6.2.0, a performance improvement was made to better balance the weights of fields that participate in matching. The automatic balancing occurs during the pre-processing phase. This helps mitigate issues that might occur when fields include many values. As a result of this update, matching operation results might be impacted. To mitigate any negative impacts, you can test and adapt your existing weights to achieve the desired matching outcome. Alternatively, you can set the `ebx.clustering.balance` property to `false` in the `ebx.properties` file to restore the legacy behavior.

Changes to third-party libraries

This release contains no updates to third-party libraries.

Closed Issues

This release contains the following closed issues:

- The `log4j.property` should be removed to ensure proper logging.
- The error message displayed does not have enough details when a metadata dataset and business table cannot be linked.
- The column size in the **Manage group** screen does not work correctly when switching between views.
- The **Merge** button is active when managing groups, even under circumstances when it should be disabled.
- It is no longer possible to resize columns in the **Merge** view.

Known Issues

This release contains the following known issues:

- The following workaround is for a known issue in EBX® (CP-27860). When child dataspace or snapshots are created before enabling the EBX Match and Merge Add-on on a data model, some operations can fail after the add-on is enabled. For instance, child dataspace merge might not work correctly, or snapshots might have errors.

Workaround: Disable the add-on before performing the dataspace merge, or creating a snapshot; enable the add-on after completing the operation:

- Open the data model, navigate to *Configuration > Add-ons*.
 - To disable the add-on, tick the box next to **Match and Merge** and select **Delete** from the **Actions** menu.
 - Publish the data model.
 - After performing the desired operations, enable the add-on by re-creating the **Match and Merge** record in the **Add-ons** table.
 - Publish the data model.
- The **Evaluate matching** service returns inconsistent results for related matching fields that have the DateTime data type.
 - It is not possible to accept or reject suspects in groups that contain a large number of records.
 - Foreign key alignment of recursive foreign keys is not possible.

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- A manual merge operation cannot be performed on more than 400 records.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 6.2.1

New features

This release contains the following enhancements:

- The **Merge** screen now supports display of custom views for tables and Business objects. This feature had previously been implemented for only the **Manage group** screen, but is now available for both screens.
- The first time users access the **Manage group** and **Merge** screens, they now display as they are configured in EBX®. If multiple views are configured for a user, they are prioritized as follows: favorite view, recommended view, then default view. Any changes users make to the view, or record orientation (vertical or horizontal) are maintained and displayed the next time users log in.
- Administrators can now specify a default view mode (horizontal or vertical) when configuring matching tables.

- The **Manage group** view now includes filter options. You can enable pre-defined filters that hide any merged records from the view, or hide all fields except those configured to participate in matching operations.
- The display of the add-on's services now reflects the permissions set on the dataset. If a service is hidden for a user on a dataset, it no longer displays as an option in the add-on's screens.
- Updates were made to improve performance when the add-on creates or modifies records.
- It is now possible to configure computed value fields to participate in matching operations.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Note:

As the add-on UI is under development, some images in the documentation might not exactly match those in the UI.

Changes in Functionality

This release contains no changes in functionality.

Upgrade notes

Take into account the following when upgrading from TIBCO EBX® Match and Merge Add-on version 3.0.0:

- When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.
- If upgrading from any version prior to 6.2.0, a performance improvement was made to better balance the weights of fields that participate in matching. The automatic balancing occurs during the pre-processing phase. This helps mitigate issues that might occur when fields

include many values. As a result of this update, matching operation results might be impacted. To mitigate any negative impacts, you can test and adapt your existing weights to achieve the desired matching outcome. Alternatively, you can set the `ebx.clustering.balance` property to `false` in the `ebx.properties` file to restore the legacy behavior.

Changes to third-party libraries

The Spring framework was updated to version 6.1.15.

Closed Issues

This release contains the following closed issues:

- When using staging to import matching policies and a search strategy does not exist in the target environment, the message should not include all the additional information related to the matching policy.
- The content of warning messages should be improved when using staging to import configuration options.
- The **Search before create** screen takes a long time to load when a decision tree has many field types.
- The survivor field value is not correct when there are multiple records that satisfy the merge function conditions.
- When performing a manual merge and fields have read-only permission, the survivor rules should be applied.
- When performing a manual merge on Business objects and fields have read-only permission, the survivor rules should be applied.
- The **Merge** button still displays after merging records and when a manual merge is no longer required.
- The loading screen gets stuck when trying to open a preview on the **Merge** screen.

Known Issues

This release contains the following known issues:

- The **Evaluate matching** service returns inconsistent results for related matching fields that have the DateTime data type.
- It is not possible to accept or reject suspects in groups that contain a large number of records.

- Foreign key alignment of recursive foreign keys is not possible.
- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- A manual merge operation cannot be performed on more than 400 records.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 6.2.0

New features

This release contains the following enhancements:

- The **Ensure group management** parameter was added to the **Manage group** workflow service. When this option is enabled, the workflow's **Accept** option only displays if the current group is resolved and has no suspect records.
- The **Back to the main view** button was re-labeled to **Review the merged group** to more closely describe the button's behavior.
- The Full text algorithm was updated to improve the consistency of matching operation results.
- When merging records or managing groups, a new vertical display option is available. This option displays the records in columns instead of rows. Administrators can set the default view with the new `ebx.addon.mame.manageGroup.record.display` property in the `ebx.properties` file. Use the values `vertically` or `horizontally` to define the view

orientation. Please note that the vertical display option is not available when merging business objects.

- The UI for viewing pre-processing information was improved as follows:
 - Now, it only displays a field configured for matching one time, even if it is used in multiple decision tree nodes.
 - A value column was added to display the values being compared.
 - The fields are grouped by their tables.
- The UI for the decision tree result explanation was updated to improve readability.
- It is now possible to transfer configuration information for data models located inside a registered module via the EBX® staging feature.
- The foreign key alignment phase was added to the progress bar when running a matching operation.
- A warning is now displayed if an algorithm that is incompatible with the defined search strategy is selected in the decision tree. This warning no longer displays when evaluating matching results.

This release contains the following performance, technical, and API enhancements:

- The pre-processing phase now takes into consideration the decision tree content to reduce the cluster size, which can lead to improved performance.
- During the pre-processing phase the balancing of matching field scores was improved so that each field equally affects the final similarity score.
- When simulating a match operation using the API, a procedure is no longer opened every time. This prevents a lock of the dataspace and any negative effects on `TableTrigger`.
- The separate procedures for matching, grouping, and merging were merged into a single procedure.
- Performance was improved when running a match operation with a selection of records.
- The REST API to accept or reject suspect records is now available.
- It is now possible to create a primary key generator to ensure that specific values from a composite primary key are survived into a golden record.

Note:

As the add-on UI is under development, some images in the documentation might not exactly match those in the UI.

Changes in Functionality

This release contains no changes in functionality.

Upgrade notes

Take into account the following when upgrading from TIBCO EBX® Match and Merge Add-on version 3.0.0:

- When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.
- If upgrading from any version prior to 6.2.0, a performance improvement was made to better balance the weights of fields that participate in matching. The automatic balancing occurs during the pre-processing phase. This helps mitigate issues that might occur when fields include many values. As a result of this update, matching operation results might be impacted. To mitigate any negative impacts, you can test and adapt your existing weights to achieve the desired matching outcome. Alternatively, you can set the `ebx.clustering.balance` property to `false` in the `ebx.properties` file to restore the legacy behavior.

Changes to third-party libraries

This release contains no changes to third-party libraries.

Closed Issues

This release contains the following closed issues:

- The record's state and relationship are not rolled back when cancelling a matching operation after the exclusion phase.
- A warning message is not included in the validation report when the default search strategy is used with the fuzzy algorithm.
- A warning message is not included in the validation report when the phonetic search strategy is used with the distance algorithm.

- Locked fields are not fixed in place when scrolling in the **Merge** screen.

Known Issues

This release contains the following known issues:

- The **Evaluate matching** service returns inconsistent results for related matching fields that have the DateTime data type.
- It is not possible to accept or reject suspects in groups that contain a large number of records.
- Foreign key alignment of recursive foreign keys is not possible.
- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- A manual merge operation cannot be performed on more than 400 records.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 6.1.3

New features

This release contains the following updates focused on performance improvements:

- Foreign Key alignment during the merging phase was optimized.
- Instead of relying on snapshots, a new method is used to check the differences between transactions.

- The API was updated to improve performance when getting lists of suspects or records contained in a group.

This release contains the following enhancements:

- New rules are applied to survivor record selection for the instances where no records, or multiple records satisfy required conditions.
- When running the **Purge old results** service, out-of-date matching selections are now also purged.
- Foreign key alignment of composite primary keys during automatic merge was improved.
- The relationship between selected records is now removed for any suspect record that is accepted as a match when managing groups.
- The readability of information written to the logs regarding foreign key alignment was improved.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains the following changes to functionality and behavior.

- If a field configured to participate in matching meets either of the following criteria, it is automatically disabled during the pre-processing phase to improve performance:
 - The field's type is `boolean`.
 - The entire column is `null`.
- Special characters are now preserved when using the add-on's distance matching algorithms.

Changes to third-party libraries

The Spring framework was updated to version 5.3.29.

Closed Issues

This release contains the following closed issues:

- The **Merge** and **Purge old results** legacy services still display on the **Merged Record Association** screen.
- Field values that satisfy the merge policy requirements are still preselected when the survivor record is read-only.
- A `NullPointerException` occurs in log4J2 during Liberty startup.
- The incorrect survivor field value is used when multiple records satisfy merge function conditions.
- When auto-merging, do not select the smallest primary key to be the survivor record when no records satisfy the **Record selection** function.
- The inline matching is marked as **Deprecated** in the breadcrumb when running the service.
- For Auto-merge, result of aligning some foreign and primary keys are incorrect if the keys are set to read-only on dataset permissions.
- For Auto-merge, alignment of foreign and primary keys is incorrect when a related table has a primary or foreign key, and auto-incremented primary key, and one of those keys is read-only by dataset permission.
- For Auto-merge, alignment of foreign keys does not work for all related tables if one FK field from any related tables is set to Read-only.
- For Auto-merge, alignment of foreign and primary keys is incorrect when a related table has a primary or foreign key, and auto-incremented primary key, and one of those keys is read-only by Access rule.
- The align foreign key function (in another dataspace) does not work correctly when the merge service runs in a child of the referenced dataspace configured on the data model.
- Medium and large numbers of merged records cannot be modified (changed to <unset>).
- The selected auto-Golden records in Related tables incorrect after merged several auto-created records by selecting each record.
- The result column in the decision tree explanation popup disappears when using an evaluation method which other than **Average**.
- A blank **Merge related** screen when creating a new golden on a business object.

Known Issues

This release contains the following known issues:

- The **Evaluate matching** service returns inconsistent results for related matching fields that have the DateTime data type.
- It is not possible to accept or reject suspects in groups that contain a large number of records.
- Foreign key alignment of recursive foreign keys is not possible.
- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- A manual merge operation cannot be performed on more than 400 records.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 6.1.2

New features

This release contains the following new features:

- It is now possible to use the replication feature on Business objects.
- Data comparison nodes now have the option to use operators or algorithms to determine how to values are compared.
- The following user experience enhancements were made:

- The display order of fields in merge policies was updated so that **Auto create new golden** is now at the top. This is more intuitive as this value can impact the availability of other options.
- The explanation for the evaluation of matching execution now includes additional parameters to better explain the results.
- Performance was improved when changing the state of selected records to .

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes to functionality and behavior.

Changes to third-party libraries

The FasterXML jackson-databind library was updated to version 2.15.2.

Closed Issues

This release contains the following closed issues:

- An exception is thrown when no duplicates are found in a group.
- All records display in the **Merged** association tab when opening a record that is not golden.
- A different method should be used to detect whether a table belongs to a Business object, or individual table configuration.
- After being deleted in the **Matching fields** tab, fields in a data comparison node cannot be re-created that use a range data type (Date, Time, Date and time, integer, decimal).
- When using staging to import the matching configuration, it fails if the decision tree in the source environment uses **Predicate**.

- The decision tree result is incorrect when: 2 or more matching fields are used, one field is from a related table, and its values are null due to being excluded by the **Exclude related records where** setting.

Known Issues

This release contains the following known issues:

- Foreign key alignment of recursive foreign keys is not possible.
- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- A manual merge operation cannot be performed on more than 400 records.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 6.1.1

New features

This release contains the following new features:

- The add-on was adapted to ensure compatibility with the EBX® staging feature.
- A page loading indicator now displays when performing a manual merge.

- When configuring inline matching for a workflow, it is now possible to use the **Matching policy** parameter to specify a matching policy other than the default policy configured for the table. The default policy is used when this parameter is undefined.
- All snapshots related to the add-on now include a `MAME_` prefix.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains the following changes to functionality and behavior:

- The **Search before create** service no longer needs to be enabled by an administrator. Additionally, when running the service, the matching policy is selectable. The chosen matching policy will determine which fields are available to search and the criteria used to locate potential matches. The **Search before create** service is still available from the matching policy configuration, however it is marked as deprecated.
- When a matching table configuration is deleted, now all related replications are also removed.
- The following default options were changed in comparison node configurations:
 - The default value for **Comparison function** is now **All fields match**.
 - For the Full text, Fuzzy, and Hybrid fuzzy algorithms, the **Minimum score** now defaults to a value of 100.
- The **Consolidated view** was deprecated. It is still available to access, but it is planned to be removed in a future release.

Changes to third-party libraries

The Spring framework was updated to version 5.2.24.

Closed Issues

This release contains the following closed issues:

- The screen does not display correctly in a data workflow when the matching field of the default matching policy is deleted.
- Color properties are not applied correctly.
- The tooltip for 'Enable Matching' is not localized in the Administration pane.
- Records that include the '+' character in the primary key cannot be viewed when accessing their group.
- Some records are incorrectly grouped after running a matching operation.
- The undo option is still clickable for values where it should not be available.
- The Manage group screen does not display correctly after changing the view.
- Records' false positive icons are not removed in the Manage group screen after deleting the records from the metadata service.
- The **Search before create** service cannot be executed when the **Matching fields** tab includes invalid matching fields.
- In workflows, the **Search before create** screen cannot be re-accessed after performing an inline merge and updating a record in list of potential duplicates.
- The add-on's manage group feature cannot be accessed when primary keys include a foreign key.
- Redundant rows can display in the **Manage group** screen.
- A foreign key field can be configured to use other algorithms besides the Exact algorithm.
- The error message from previous phase is not cleared after changing field.
- Nothing happens when clicking the Preview button on primary key or foreign key field in the Lineage screen.
- Display ALL services under Actions menu in the association of Merged records.
- The screen does not display correctly after previewing a record in the Manage group screen.
- Matching results are provided using the light relationship instead of the strong relationship when both are configured on a business object.

Known Issues

This release contains the following known issues:

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.

- A manual merge operation cannot be performed on more than 400 records.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 6.1.0

Released: March 2023

New features

This release contains the following new features:

- The add-on was adapted to ensure compatibility with the way EBX® now handles URL encoding. See the EBX® release notes for additional information.
- Group management:
 - The **Manage group** screen user interface and user experience were updated and improved.
 - Any existing custom table views can now be applied in the **Manage group** screen.
 - All group management functions are now available when accessing this feature through a workflow. Additionally, the **Close** button is now available on this user service in workflows.
- Manual merge:
 - The user interface and user experience were updated and improved in this release for the **Merge** view. The view includes a new merge progress tracker that includes sub-steps when merging a business object's child tables.

- A new **Undo** button was added to the **Merge** view. This button allows you to revert the most recent values selected for a golden record. The reverted value is highlighted in the UI.
- The merge **Summary** screen was updated to display successfully merged records along with the icons to access the record's detailed view.
- Hybrid fuzzy algorithm:
 - When using the hybrid fuzzy algorithm in a decision tree node, you can now customize the values used by the algorithm for: Synonyms, Levenshtein 1, and Levenshtein 2.
 - The Levenshtein transformation logic was updated to ensure that short strings return more accurate results.
- The **No fields match** option was removed from decision tree comparison nodes. If you used this option in one of your add-on configurations, navigate to: *Administration > TIBCO EBX® Match and Merge Add-on > Table activation and settings*. The **Status** column will alert you to any configurations where you must select a different comparison function.
- The option to automatically trigger matching operations is now available for business objects. Only the **After submission** option is available to use with business objects, either **On creation** or **On update**.
- Evaluation of matching:
 - When evaluating matching results and viewing results for a data comparison node, the displayed information now includes the table that each compared field belongs to.
 - The similarity results were updated to include additional information of why records do not complete the pre-processing stage.
- New Java APIs were introduced to:
 - Evaluate matching results.
 - Retrieve the trusted sources list.
 - Perform multiple merges in one transaction.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

The behavior of the **Modify merged records** option was reverted. You can once again choose what the add-on does when users modify merged records. An additional option was also included that is a combination of **Keep records in group** and **Change to <unset>**.

Changes to third-party libraries

The SnakeYAML library was updated to version 2.0.

Closed Issues

This release contains the following closed issues:

- An inconsistent matching result is shown in the diagram when evaluating matching results.
- An error displays after running the **Evaluate matching** service multiple times.
- Matching cannot be evaluated when using a group's field and a related field as matching fields.
- When managing relationships during a manual merge, it takes a long time to display records.
- There is no new golden record but the **Apply merge policy** button is still enabled.
- An incorrect error message displays when manually merging and the main table's primary key is not auto-incremented.
- The record order in when managing a group in the **Consolidated view** is incorrect.
- An incorrect tool-tip displays after modifying values in a golden record.
- The matching field is changed when duplicating a matching policy.
- Hidden field values are not updated after the merge process finishes.
- The warning banner does not display when accessing a group that includes at least one hidden record.
- Records are not sorted correctly when accessing the **Merge** view when managing a group.
- When hovering over the **False positive** icon in the **Manage group** screen, the cursor displays as a hand.
- A false positive relationship is not created after merging records in the inline screen.
- The **State** is duplicated the **Identified as false positives** tab.
- Incorrect matching results are returned when using light relationships (FK from matching table to target table OR over a link table).
- The back button does not work as expected after completing a merge.

- Foreign keys do not display correctly in a business object's **Manage group** view.
- Some text is not translated into French.
- All add-on services still display in the **Actions** menu after pausing the corresponding matching configuration.
- Matching cannot be evaluated for selected records when using a matching policy without having a merge policy defined in the related table.
- The decision tree result is inconsistent when evaluating matching for a URI field through a relationship.
- An error displays after unmerging an auto-created golden record that had been merged multiple times.
- Unexpected information is included in the log file.

Known Issues

This release contains the following known issues:

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.6.1

Released: February 2023

New features

This release was updated to ensure compatibility with TIBCO EBX® Add-ons Bundle version 5.6.1.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

This release contains no changes to third-party libraries.

Closed Issues

This release includes no closed issues.

Known Issues

This release contains the following known issues:

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data

Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.

- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.6.0

Released: November 2022

New features

This release contains the following new features:

- The UI was updated for manual merge operations on the tables that comprise business objects.
- The stewardship UI was updated so that lineage and metadata can be viewed for a business object's related tables.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

The FasterXML jackson-databind library was updated to version 2.13.4.2.

Closed Issues

This release includes the following closed issues:

- The manual merge view is slow to load if the table has a merged records association.

- The **Align foreign keys** step does not display when a merge policy attribute is defined and **Use for manual merge** is set to **No**.
- Hidden field values are still highlighted when selecting a survivor field by merge function.
- Matching state is cascaded to child levels when unmerging even though the child records are excluded.
- When using matching fields from related tables with the Exact algorithm and **Both value are null = Match**, there are inconsistent matching result between Evaluate Matching and See details.
- An HTML block is displayed instead rendering the Decision tree and Business object activation.
- Deleted records display in the **Data selection** view of the **Supersede survivorship** service.
- The **Manage group** screen should be reloaded after a successful manual merge operation.

Known Issues

This release contains the following known issues:

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.5.0

Released: September 2022

New features

This release contains the following new features:

- Matching and merge operations can be configured on multiple tables that comprise EBX® business objects.
- The **Manage group** view is able to display records from multiple tables in a business object.
- To improve the user experience, the progress bar was updated.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

The Apache Taglibs library was removed from this release.

Closed Issues

This release includes the following closed issues:

- The content does not display correctly when there are more than 50 errors in a validation report on the **Consolidated view**.
- There is a disparity between matching results when using **Evaluate matching** vs **Run match**.
- After rejecting a suspect record, users are navigated away from the Manage group screen and back to the tabular view.
- `DecisionTreeComparator` is not thread-safe.

Known Issues

This release contains the following known issues:

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.4.1

Released: July2022

New features

This release contains no new features.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

This release contains no changes to third-party libraries.

Closed Issues

The merge process is slow when changing the record state to `Unset`.

Known Issues

This release contains the following known issues:

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.4.0

Released: June 2022

New features

This release contains the following new features:

- The following algorithms were added:

- Phonetic full text: An algorithm best used for strings. It can recognize the phonetic equivalent of two words spelled differently using the Beider-morse phonetic tokens. The comparator takes into account the synonyms and stop words defined in the data model.
- Hybrid fuzzy: An algorithm best used for string and text data types. It can capture complex relationships between two field values by automatically customizing transformations for a given criterion. Points of note about this algorithm:
 - Does not return a 100% matching score on synonyms.
 - Strings that are similar, but contain characters like spaces or dashes do not return a 0% matching score.
 - When strings fall within the acceptable distance score they are not considered as the exact same. For example, a distance of 1 returns .95, a distance of 2 returns .9, etc.
- A script task was added to run a matching operation from a workflow.
- Fields can now have a weight of 0. Setting the weight to 0 excludes the fields from the pre-processing phase.
- An API was added to check whether the add-on has any ongoing background operations.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

The options to specify what happens after a merged record is modified were removed. Now, the following behavior is applied when a record is in the merged state and the modified field:

- Is *not* configured to use in matching operations: The add-on executes an automatic merge using the settings defined in the default matching policy's merge policy.
- Is configured to use in matching operations: The record is changed to the <unset> state.

Changes to third-party libraries

This release contains no changes to third-party libraries.

Closed Issues

This release contains the following closed issues:

- Changing record state to **Unset** fails when many transactions are committed in a short period of time.
- The **Evaluate matching** and **Run match** services return inconsistent results.
- There is conflict between the **Evaluate matching** and **See details** results.
- Results cannot display after running **Evaluate matching** on a display with 1920X1080 resolution at 150% zoom.

Known Issues

This release contains the following known issues:

- Linked records are duplicated in the link table when one of the foreign keys is not part of the primary key.
- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.3.5

Released: June 2022

New features

This release contains no new features.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

This release contains no changes to third-party libraries.

Closed Issues

A performance issue occurs during the grouping phase when running matching on an entire table.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.

- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.3.4

Released: May 2022

New features

This release contains no new features.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

The Spring framework was updated to version 5.2.22.

Closed Issues

This release contains no closed issues.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.

- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.3.3

Released: April 2022

New features

This release contains no new features.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

This release contains no changes to third-party libraries.

Closed Issues

This release contains the following closed issues:

- The **Merge view** and the **Lineage view** cannot be accessed when all fields in a group or table are computed fields.
- Performance is slow when fetching values n to n relationships.
- When running the **Evaluate matching** service the field values do not display correctly when using the same field identities, but in different tables.
- Boolean fields cannot be merged when **Check null input** is enabled.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.3.2

Released: April 2022

New features

This release contains no new features.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

The Spring framework was updated to version 5.2.20.

Closed Issues

This release contains no closed issues.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.3.1

Released: March 2022

New features

Performance was improved to retrieve groups containing suspect records.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no changes in functionality.

Changes to third-party libraries

This release contains the following changes to third-party libraries:

- The Spring framework was updated to version 5.2.19.
- The FasterXML/jackson-databind library was updated to version 2.13.2.1.

Closed Issues

This release contains no closed issues.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you

can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.

- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.3.0

Released: February 2022

New features

This release includes the following new features:

- You can now replicate matching metadata. This allows you to use third-party tools to access and query the metadata.
- The following algorithms are now available:
 - Fuzzy full text: This algorithm works best for strings and can find the similar and fuzzy matches of the words in the compared values. It is based on the Levenshtein distance algorithm. The comparator takes into account the synonyms and stop words defined in the data model.
 - Range: This algorithm matches values within a predefined range. The two values are considered a match if the distance between two values is within the specified range.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains the following changes in functionality:

- When **Auto create new golden** is enabled, the requirement was removed for the matching table's primary key to be an auto-incremented integer.

- When executing add-on operations, only the add-on's table triggers are disabled. All user custom triggers remain enabled. Note that table triggers related to metadata linked fields are not recommended as it might affect the accuracy of match and merge operations.

Changes to third-party libraries

This release contains no changes to third party libraries.

Closed Issues

When matching is set to **After submission** and batch creation of records occurs using the EBX® API, the number of submitted or created records is limited to 2000 records in an API call for tables that have a single primary key. For a composite primary the limit is divided by the number of fields that makes up the primary key. For example, with a table that has a primary key comprised of 2 fields, the limit would be 1000 records. When your requirements exceed these limits, it is recommended that you pause the matching trigger, load the data, and resume matching. Also, note that when matching is triggered the dataspace is locked until matching completes. Any incoming requests during this time are blocked and will return errors.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- For foreign key alignment, the FKs that are multi-value fields cannot be aligned (either manually or automatically).

Version 3.2.1

Released: December 2021

New features

This release includes no new features.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

The Apache Log4j library was upgraded to version 2.17.1.

Changes to third-party libraries

This release contains no changes to third party libraries.

Closed Issues

This release contains no closed issues.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you

can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.

- When matching is set to **After submission** and batch creation of records occurs using the EBX® API, the number of submitted or created records is limited to 2000 records in an API call for tables that have a single primary key. For a composite primary the limit is divided by the number of fields that makes up the primary key. For example, with a table that has a primary key comprised of 2 fields, the limit would be 1000 records. When your requirements exceed these limits, it is recommended that you pause the matching trigger, load the data, and resume matching. Also, note that when matching is triggered the dataspace is locked until matching completes. Any incoming requests during this time are blocked and will return errors.

Version 3.2.0

Released: November 2021

New features

This release includes the following new features:

- You can now exclude records from matching operations by specifying a field value. If a record contains the specified value, it is not included in matching operations. Using this functionality can help improve performance and improve the quality of matching results.
- You can now modify records that were included in a merge operation.
- The new **Evaluate matching** service allows you to test matching results. This can help you fine tune configuration settings to achieve desired results.
- Adding an association to your data model allows you to view: a hierarchy of golden and merged records, and a tab when viewing golden record details that displays all merged records that target the golden.
- Matching now supports use of non-Latin based text.
- A new merge function on field survivorship allows you to add a constant value that is merged into golden records on survivorship.
- The Full Text algorithm is now available to use in matching for fields with Text and String data types.
- You can now use a foreign key field as a field in a decision tree's comparison node.

Attention:

When upgrading from version 3.0.0, you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains the following functionality changes:

- The default mode for **Align foreign keys** is now set to **Manually**.
- The **Case sensitivity** option is now only available for fields with String and Text data types.

Changes to third-party libraries

This release contains no changes to third party libraries.

Closed Issues

This release contains no closed issues.

Known Issues

This release contains the following known issues:

- When matching using the foreign key option, you can only specify a single hop foreign key. In other words, it cannot be a foreign key to a foreign key.
- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.

- When matching is set to **After submission** and batch creation of records occurs using the EBX® API, the number of submitted or created records is limited to 2000 records in an API call for tables that have a single primary key. For a composite primary the limit is divided by the number of fields that makes up the primary key. For example, with a table that has a primary key comprised of 2 fields, the limit would be 1000 records. When your requirements exceed these limits, it is recommended that you pause the matching trigger, load the data, and resume matching. Also, note that when matching is triggered the dataspace is locked until matching completes. Any incoming requests during this time are blocked and will return errors.

Version 3.1.1

Released: September 2021

New features

This release includes updates and enhancements to improve the performance of matching operations.

Attention:

When upgrading from version 3.0.0, 3.1.0, or 3.1.1 you must update existing matching configurations. After updating configuration settings, you should re-execute matching to ensure that matching related metadata is up to date.

Changes in Functionality

This release contains no functionality changes.

Changes to third-party libraries

This release contains no changes to third party libraries.

Closed Issues

This release contains no closed issues.

Known Issues

This release contains the following known issues:

- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- When matching is set to **After submission** and batch creation of records occurs using the EBX® API, the number of submitted or created records is limited to 2000 records in an API call for tables that have a single primary key. For a composite primary the limit is divided by the number of fields that makes up the primary key. For example, with a table that has a primary key comprised of 2 fields, the limit would be 1000 records. When your requirements exceed these limits, it is recommended that you pause the matching trigger, load the data, and resume matching. Also, note that when matching is triggered the dataspace is locked until matching completes. Any incoming requests during this time are blocked and will return errors.

Version 3.1.0

Released: July 2021

New features

This release includes the following new features:

- A new REST API is available to perform matching related operations. You can use Swagger to view the API documentation. See [REST overview](#) for more information.
- Matching operations after record submission are now available for records created using the EBX® API. There are some limitations to this functionality. See the *Known issues* section below for details.
- Several matching services are now available to use in workflows and perspectives. See [Perspectives and workflows](#) for more information.

- The add-on now supports link tables when matching with relationships. Additionally, to avoid confusion with EBX®'s concept of associations, this feature was renamed to **Use relationships**.
- The add-on's matching algorithm library now includes the Soundex algorithm. This algorithm works well for short strings such as names. See [Matching algorithms](#) for more information.
- The list of algorithms is now filtered based on the search strategy configured for a matching field. Additionally, the validation report will display a warning if the configured algorithm is not compatible with the field's search strategy. See [Matching algorithms](#) for more information.
- Administrators can now set a matching policy as the default policy from the **Matching policies** tab in the **Actions** menu.
- A merge policy's record selection mode now includes the **First acquired** option to select the oldest record during survivorship.
- The **Customize source value for new golden** now supports foreign keys.
- Matching metadata can now be added to tables by creating a custom view. See [Displaying matching metadata](#) for more details.
- The add-on's **Consolidated view** the **Actions** menu where you can execute matching related operations.
- When viewing table data, you can now access the **Metadata** view's **Lineage** tab by selecting **View lineage**.
- When a group contains only one record that does not need to be merged, a new **Set as golden** button allows you to set its state to golden.

Changes in Functionality

This release contains the following functionality changes: When using the **Run match** service to run a match on all records in a table, you must now select all records in the table instead of selecting nothing.

Changes to third-party libraries

This release contains no changes to third party libraries.

Closed Issues

This release contains no closed issues.

Known Issues

This release contains the following known issues:

- The **Validation** service in the **Consolidated** view runs as expected when first executed, but slow on subsequent executions.
- In order to accurately track record lineage, table history must remain activated over time. If the table's activation status repeatedly changes, the lineage data will be inconsistent.
- Severe errors occur when multiple datasets based on the same model exist and a table from one of these datasets is activated in an add-on matching configuration. As a workaround, you can create another data model publication using the **Manage publications** service in the Data Model Assistant. After creating a dataset for the new publication, you can follow the normal procedures to configure a matching policy for tables in the new dataset.
- When matching is set to **After submission** and batch creation of records occurs using the EBX® API, the number of submitted or created records is limited to 2000 records in an API call for tables that have a single primary key. For a composite primary the limit is divided by the number of fields that makes up the primary key. For example, with a table that has a primary key comprised of 2 fields, the limit would be 1000 records. When your requirements exceed these limits, it is recommended that you pause the matching trigger, load the data, and resume matching. Also, note that when matching is triggered the dataspace is locked until matching completes. Any incoming requests during this time are blocked and will return errors.

Version 3.0.1

Released: March 2021

Bug Fixes

A documentation issue was corrected.

Version 3.0.0

Released: March 2021

Release overview

The TIBCO EBX® Match and Merge Add-on locates and merges duplicate data values. The goal is to obtain a singular record that most accurately defines a business entity. The add-on identifies these as *golden records*.

Features available to administrators include the ability to:

- Make data available to the add-on by registering and activating matching on tables.
- Configure a matching policy that determines how the add-on compares data. A matching policy also specifies whether users must start the matching process manually or if it starts automatically when data is created or updated.
- Create a merge policy that establishes how duplicate data values are merged into a golden record. Merge policies can also execute automatically after a matching process runs.
- Set up trusted sources that define a hierarchy of data sources. The add-on uses the source's trust level to during matching to choose the most accurate record and during a merge to select values to merge.

Administrators can use [Administrator Guide](#) to learn more about configuration options available in this release.

Features available to business users include the ability to:

- Select data and initiate matching and merge operations. Search for records that might match when creating a new record.
- Manage records that grouped together after running a matching process. The add-on groups records that are suspected of containing duplicate data or are positively identified as being duplicate. When managing groups you can also specify that a record is a false-positive match so that it no longer matches with the records in its group.
- View record metadata applied by the add-on to categorize records. Two different views allow users to view individual record, or an entire table's metadata. Some information displayed includes the record's history, lineage, state, and records identified as false positive matches.

Business users can learn more about available actions and view in the [User Guide](#).