

DISCOVERY HUB 18.10

RELEASE DOCUMENTATION



Contents

Introduction	3
New Features	4
Operational Data Exchange (ODX) with support for Azure Data Lake	4
Azure SQL Database Managed Instance	4
Shared Semantic Layer	4
Row Level Security	4
Hierarchies	5
Hide Settings for Unused Endpoint Types	5
Validate Model Against Endpoints	5
Data Formats	5
Display Folders.....	5
Prevent Deployment of an Endpoint in Specific Environments.....	5
Guarding	5
Prioritize Models on Execution	6
Improved Method for Assigning Semantic Objects to Endpoints.....	6
Improved Mapping of Tables	6
Added Table and Schema Reference to Tables.....	6
Consistent Dragging Direction when Creating Relations	6
Improved Display of Relations	6
Clone field	6
Consistent Read on Dynamics NAV Adapter	6
New Features in Beta	7
Save time: Keep Lookup Values Up-to-Date on Incrementally Loaded Tables	7
Improved Features	8
Improved Oracle Data Source	8
Tooltip	9
Conditional Lookup Performance and Usability Improvements.....	9
Minor General Improvements	15
Update Connection Manager	15
Deprecated Features.....	15
Qlik Modeler	15



Introduction

Dear reader, a new Discovery Hub® release has arrived!

We're excited to announce that Discovery Hub® now offers support for **Azure Data Lake**. In fact, we have completely rebuilt one of the components of Discovery Hub: the operational data exchange (ODX). The new ODX enables users to fully automate the process of flowing data in and out of an Azure Data Lake.

Are you planning to upgrade your SQL database to the cloud? Then we have some more good news! Microsoft has announced to make Azure SQL Database Managed Instance general available on October 1st. Since Discovery Hub® version 18.6 users can choose **Azure SQL Database Managed Instance** as the storage. This means that you can move to the cloud with a click of a button, the powerful automation engine behind Discovery Hub® does the rest.

One of the other objectives for this release is to dig deeper into the **shared semantic layer** and add those features that you just cannot live without. Among the new features, you will find support for Row Level Security for both SQL Server Analysis Services (SSAS) Tabular and Qlik. This means that with a single setup, you can deploy the same setup of row-level security to both endpoints. This gives even more reason to use the Shared Semantic Layer when you have multiple front-end tools in your organization.

The fourth objective for this release is to revisit one of the most commonly used features of TimeXtender Discovery Hub® - **Lookups!** We have taken a closer look at all suggestions and feature requests regarding Lookups and have combined them with our own ideas. The result is a more performant, easier to use lookup feature. I would also like to use this opportunity to thank you all for participating in our technical forum on <http://support.timextender.com>! Though we might not always respond to individual suggestions, we read, value and consider all your input!

We also have a very interesting feature in Beta for this release! With a single checkmark, you can keep your lookup values up-to-date, so you no longer have to script updates to valid tables or perform full loads to get lookup values updated. It does not get any simpler or more automated than this!

We hope you will enjoy Discovery Hub® 18.10!

Thomas Lørup Duun

Chief Technology Officer



New Features

Operational Data Exchange (ODX) with support for Azure Data Lake

The new ODX enables users to fully automate the process of flowing data in and out of an **Azure Data Lake**. This function permits organizations to populate the Azure Data Lake with virtually any number and any type of data sources available on the market. Data will be available in the same structure and quality as in the source systems but will now be accessible from one location. This takes out the pain of data preparation while also avoiding performance and security issues on source systems. In addition, companies will have the capability to connect **AI, predictive or self-service analytics tools**; enabling IT to use Discovery Hub® as a single platform for offloading all their data for business users to access and leverage.

These data loads are fully configurable enabling users to schedule loads at any specific time and frequency. Of course, the ODX fully supports incremental loading with the option to load only new and modified data from your sources and into your Azure Data Lake. Per default, the ODX synchronizes meta data from the data sources before each load to make sure you always have all data available in the ODX.

Azure SQL Database Managed Instance

The new release entirely supports **Azure SQL Database Managed Instance**, - general available from 1 October 2018 -, that represents fully managed SQL Server Instance hosted in Azure cloud — allowing customers to easily take advantage of new technology when released by Microsoft. In fact, Discovery Hub® 18.10 supports any of Microsoft SQL Servers, be it Microsoft Azure SQL DB Managed Instance, Single or Elastic Pool and of course the on-prem SQL database. We recommend MI, but if performance and refresh rate are less important, Single or Elastic pool are also an option.

Shared Semantic Layer

Row Level Security

In the new release of Discovery Hub, you can set up row-level security on models in the shared semantic layer. The security setup can be deployed to SSAS Tabular and Qlik Endpoints.

Controlling access to data on the row level can be a very useful addition to controlling access on the table or model level. One common example is a sales department where each salesperson needs access to data from their own sales region, but not everyone else's.

In this release, the correlation between users and their permissions are managed entirely inside Discovery Hub®. For the next feature-release, we have planned an option to load this information from tables instead to provide a more Dynamic Security Setup.



Hierarchies

In Discovery Hub® 18.10 you can order fields on semantic layer tables in hierarchies. All endpoint types support hierarchies. However, Tableau works slightly different in that a field can only be in one hierarchy. If you use the same field in two hierarchies, you will be warned that the hierarchy cannot be deployed to a Tableau endpoint.

Ordering fields in hierarchies can be a good way of making a model easier to grasp and use in PowerBI, Qlik Sense or Tableau. An example is a table with products that often have multiple levels of categories.

Hide Settings for Unused Endpoint Types

In the new release, you can choose to hide endpoints, that you don't use to simplify your models. You can select what endpoint types you want to display settings for on models in the semantic layer. These settings include, among others, data formats, data categories and display folders on fields, cardinality on relations and row-level security. Settings that can be toggled off have info-tips next to them.

The main advantage of the shared semantic layer is that it is just that – shared. The same model can be deployed on different endpoints. However, few models need to be deployed to all supported endpoints and settings that are not relevant often clutter the interface.

Validate Model Against Endpoints

Discovery Hub® 18.10 adds a validation option on models in the semantic layer. The validation detects and shows a list of issues in the model that would cause trouble on the model's endpoints. The list is split by endpoint since the different types have different requirements. Some of the things that are checked are circular references (Qlik endpoints), tables without relations (Tableau), and custom fields without data type (Tableau).

Data Formats

You can now configure data formats and categories on fields which will be applied to SSAS Tabular endpoints and used in frontends to decide how to display data. For instance, text categorized as "Web URL" can be displayed as a link.

Display Folders

"Display folders" is a feature of SSAS Tabular that allows you to organize the fields, measures and hierarchies of a semantic table by adding them to folders. This makes the model easier to browse, especially if it has a lot of fields. With display folders, you can e.g. bucket all address information on a customer table in a "Geography" display folder.

Prevent Deployment of an Endpoint in Specific Environments

In the new release, you can set an endpoint to be inactive in specific environments. This makes it easier to manage endpoints and prevent overriding production endpoints during development.

Guarding

Semantic Models can now be guarded for deployment and/or execution.



Prioritize Models on Execution

Semantic Models can now be prioritized in the execution engine, which gives the user an option to influence the execution order to get certain items finished earlier than others.

Improved Method for Assigning Semantic Objects to Endpoints

It is now much faster and easier to add Semantic objects to endpoints. The multistep approach has been reduced to a one-step method: right-click and select "specify endpoint" (works much like how you add tables to perspectives).

Improved Mapping of Tables

The methodology for mapping Data Warehouse or Business unit tables into a semantic model table has been improved and simplified. By default, no fields are selected, which is the behavior that has been requested by a lot of users. With a single click, you can select all fields.

Added Table and Schema Reference to Tables

When using tables from different levels in the project – i.e. Business Unit and Data Warehouse or from different schemas in the same level – it was not possible to identify where the individual tables originated. Thus, a Tooltip has been added that shows which database and schema the originating table belongs to.

Consistent Dragging Direction when Creating Relations

The dragging direction when creating relations between two tables in the semantic model has been reverted to be consistent with the behavior on the Data Warehouse and Business Unit Levels.

Improved Display of Relations

Relations are displayed under the table it has a relation to, instead of in one big unsorted list of relations.

Clone field

We have added a new context menu command "Clone field" on all fields, as well as a new dialog for cloning fields (incl. Conditional Lookups, Supernatural Keys, Custom Hash Fields, transformations, standard data fields). Two options appear in the dialog: Cloning structure (incl. transformations) or data (only adding transformations for cloned data).

Consistent Read on Dynamics NAV Adapter

When you extract data from a data source that is in use, you can end up with inconsistent data where a foreign key in a record points to a record in another table that does not exist.

Discovery Hub® 18.10 can use the fact that Dynamics NAV has a timestamp on every table to minimize this problem. If you enable the new consistent read setting, Discovery Hub® will only extract records older than the start of the data extraction. Though this will not help with any inconsistencies present



when the data extraction begins, but it will prevent inconsistencies caused by data added while extraction is running.

New Features in Beta

Beta features have been thoroughly tested in-house but need additional testing in production scenarios to make sure they have the required performance and behavior in all cases.

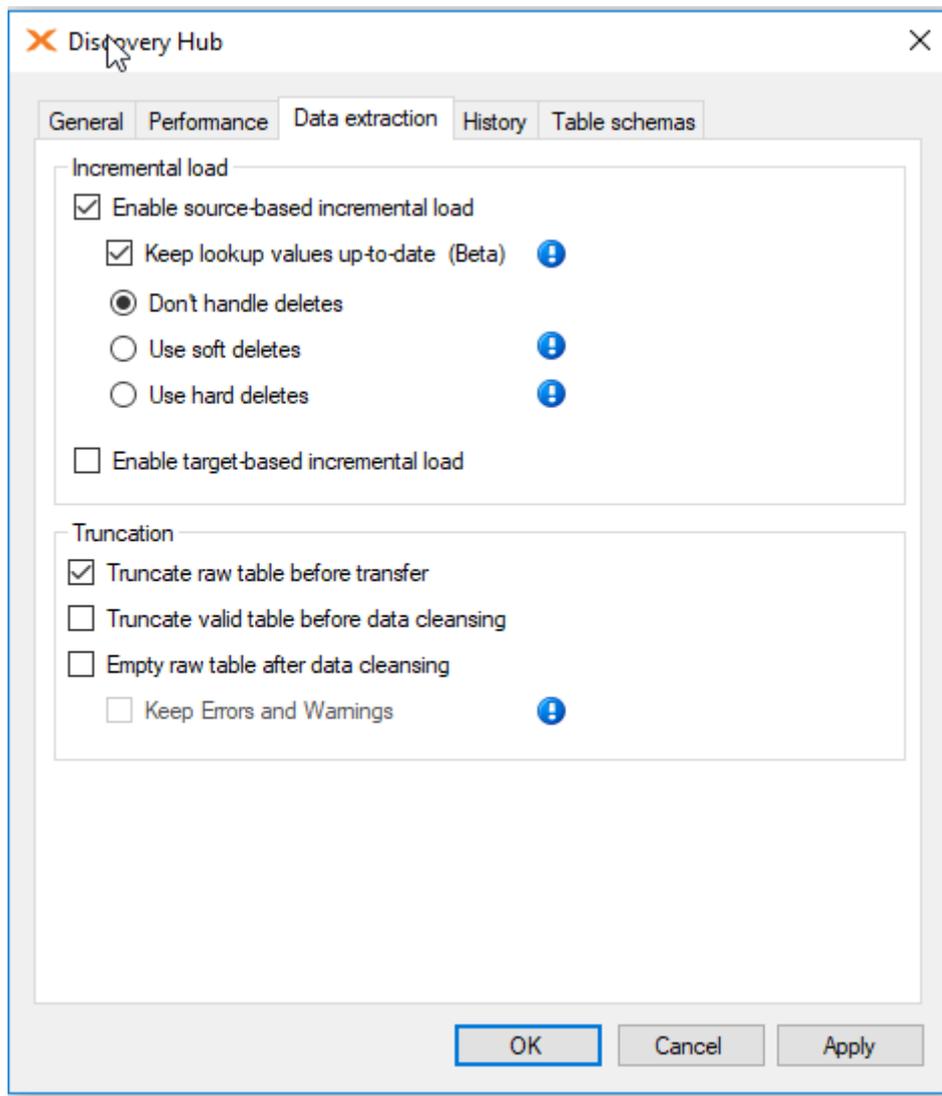
Save time: Keep Lookup Values Up-to-Date on Incrementally Loaded Tables

One of the challenges of using incremental load is that the value of a lookup field on an incrementally loaded table does not change if the value being looked up changes. The lookup field value is calculated when the row is loaded so the usual solution to this problem is to schedule regular full loads of the table.

In Discovery Hub® 18.10, we have added a new setting on incrementally loaded tables, “Keep lookup values up-to-date”. When this setting is enabled, Discovery Hub® will do comparisons to see if any values being looked up have changed and will update the lookup fields accordingly.

While the logic is highly optimized for speed, enabling the feature does come with a performance penalty in some scenarios.





Improved Features

Improved Oracle Data Source

This release contains some improvements to the Oracle data source which makes it easier to set up the connection and makes the connection Oracle-version independent.

The data source connector now supports EZConnect which means that you can simply enter the connection information (server, port and service name). In previous versions of Discovery Hub®, you had to enter an exact "net service name" corresponding to an entry in the "tnsnames.ora" configuration file.



To connect to an Oracle database, the connector needs a provider to be installed on the system. Previous versions of Discovery Hub® required a specific provider while the new version supports four different providers and will automatically use the best one available.

More Information on Tooltips

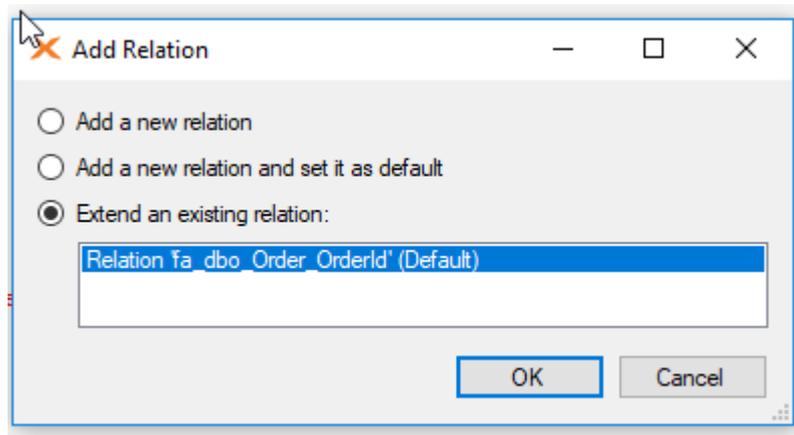
Tooltip is now on all fields with information on Transformations. Tooltip on lookup fields will now show information from the lookup data field, joins, conditions, etc. and there is improved tooltip on conditional lookup fields. The tooltips react to all relevant changes to fields and relations, always showing up-to-date information. Lastly, the tooltip doesn't disappear as fast as before (it is shown for 30 seconds).

Conditional Lookup Performance and Usability Improvements

Lookups! We have taken a closer look at all suggestions and feature requests regarding Lookups and have combined them with our own ideas.

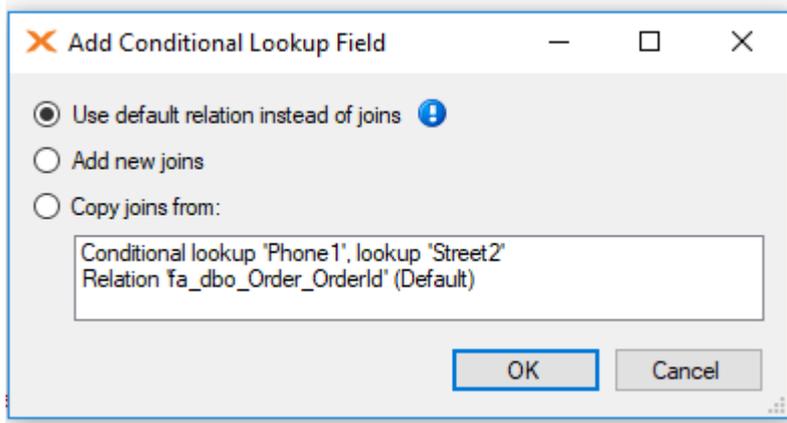
Here is a list of the improvements and fixes:

New feature: When dragging and dropping fields, a new easy to use dialog for adding relations is shown, displaying existing relations with a tooltip, possibility of extending an existing relation, or adding a new default relation, or non-default relation.

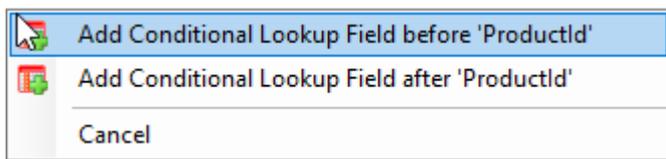


New feature: When adding a new conditional lookup field or a new lookup field, a new simplified, comprehensible dialog for adding lookup joins is shown, with a possibility of copying joins not only from relations but also from existing lookup fields. This dialog also shows details of all joins in tooltips. Only different joins from existing lookups are shown (duplicates are removed from the list).

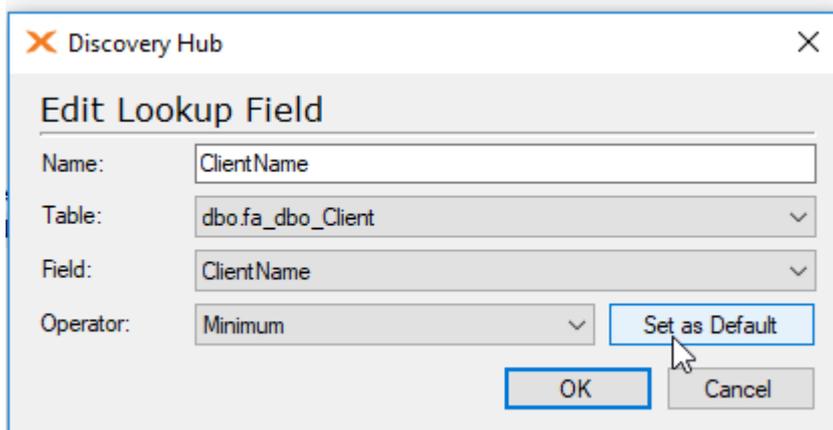




New feature: Dragging and dropping of fields in another table with the **right** mouse button creates a new conditional lookup field instead of a new relationship and allows for easy positioning of the conditional lookup field before or after the dropped field.

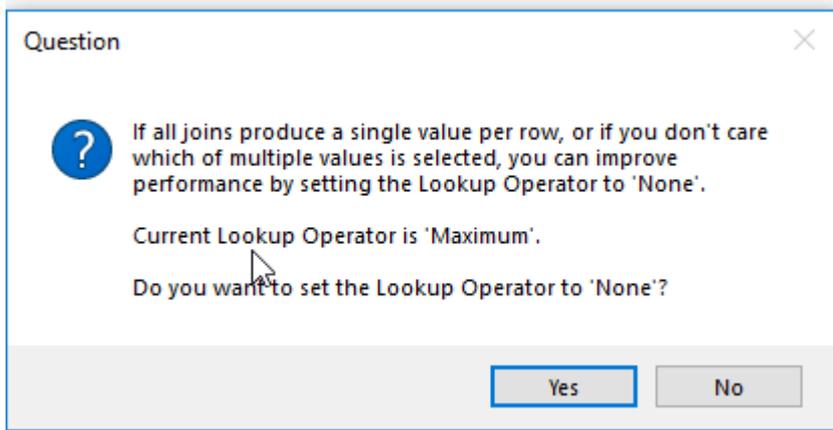
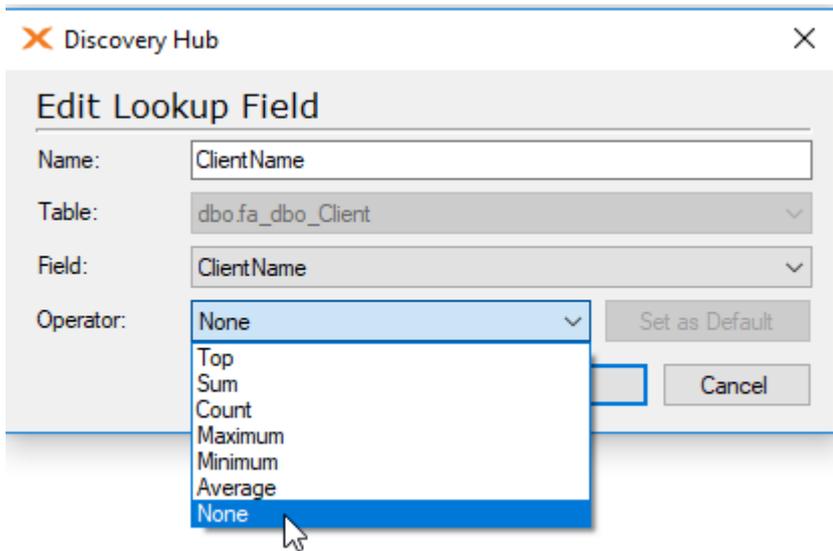


New feature: In the Add/Edit Lookup Field dialog, you can now select which Operator you want to "Set as Default".

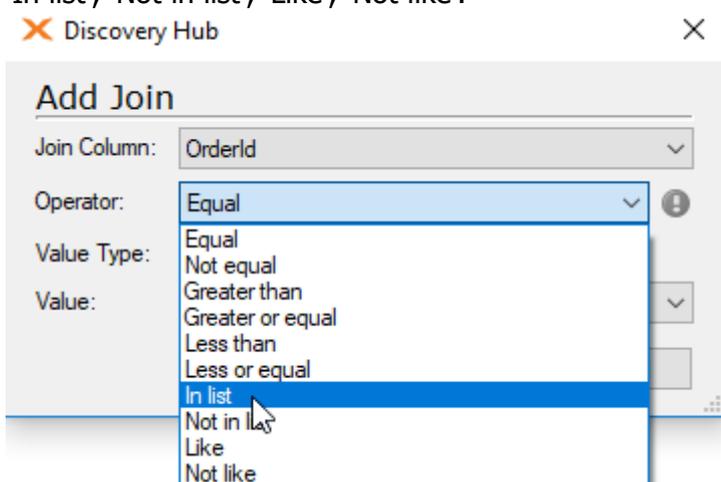


New feature + Improved performance: New lookup operator 'None' greatly improves performance in some scenarios + automatic changing (after confirmation) of the lookup operator to 'None' when adding joins different than 'Equal'.





New feature: In the "Add Join" dialog, 4 new lookup join fixed value operators have been added: 'In list', 'Not in list', 'Like', 'Not like'.



New feature: Copying Lookup Fields within the same table via drag and drop (the right mouse button also offers two options: add or replace)

New feature: Copying Lookup Joins within the same table via drag and drop (the right mouse button also offers two options: add or replace)

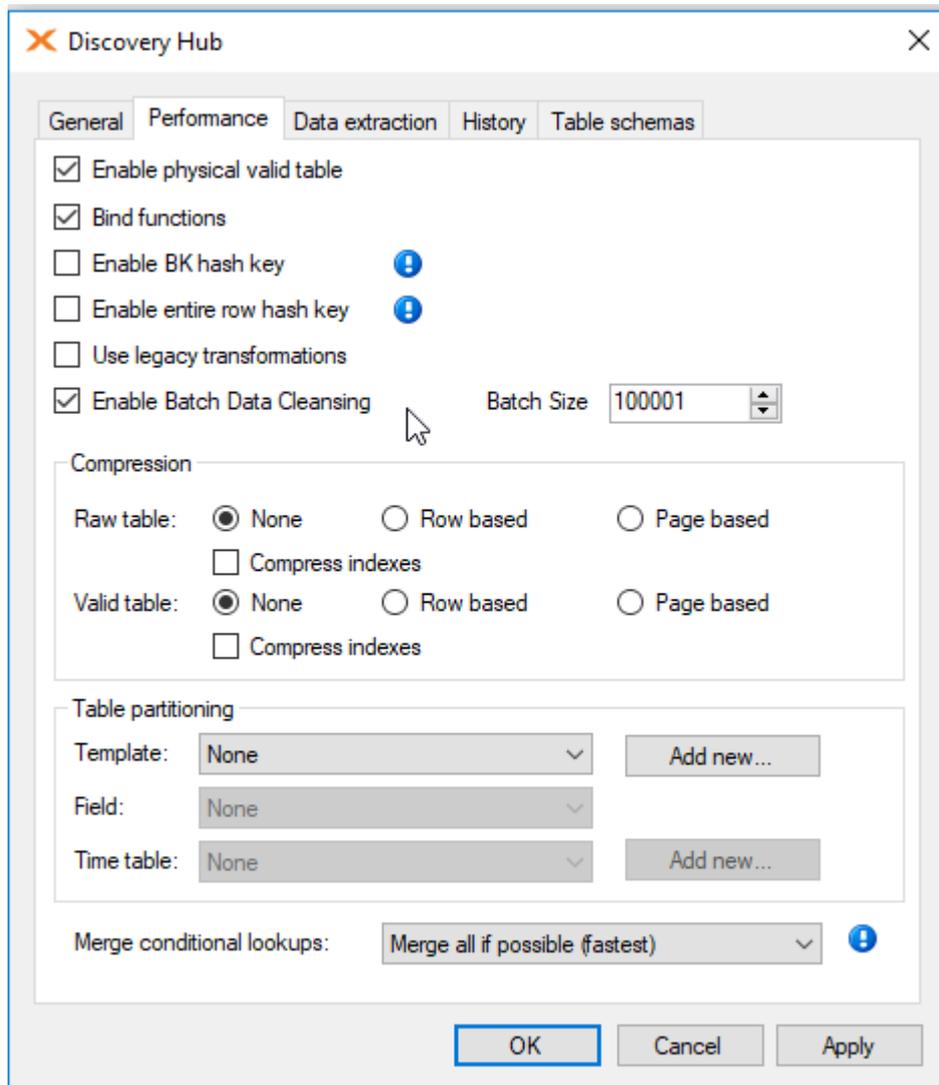
New feature: Copying Lookup Sortings within the same table via drag and drop (the right mouse button also offers two options: add or replace)

New feature: Copying Lookup Conditions within the same table via drag and drop, also between lookups and transformations, and vice versa (the right mouse button also offers two options: add or replace)

New feature: Automatic renaming (after confirmation) of lookup fields and conditional lookup fields when the lookup data field is changed

Improved performance: Batch data cleansing now works also for all conditional lookups and supernatural keys

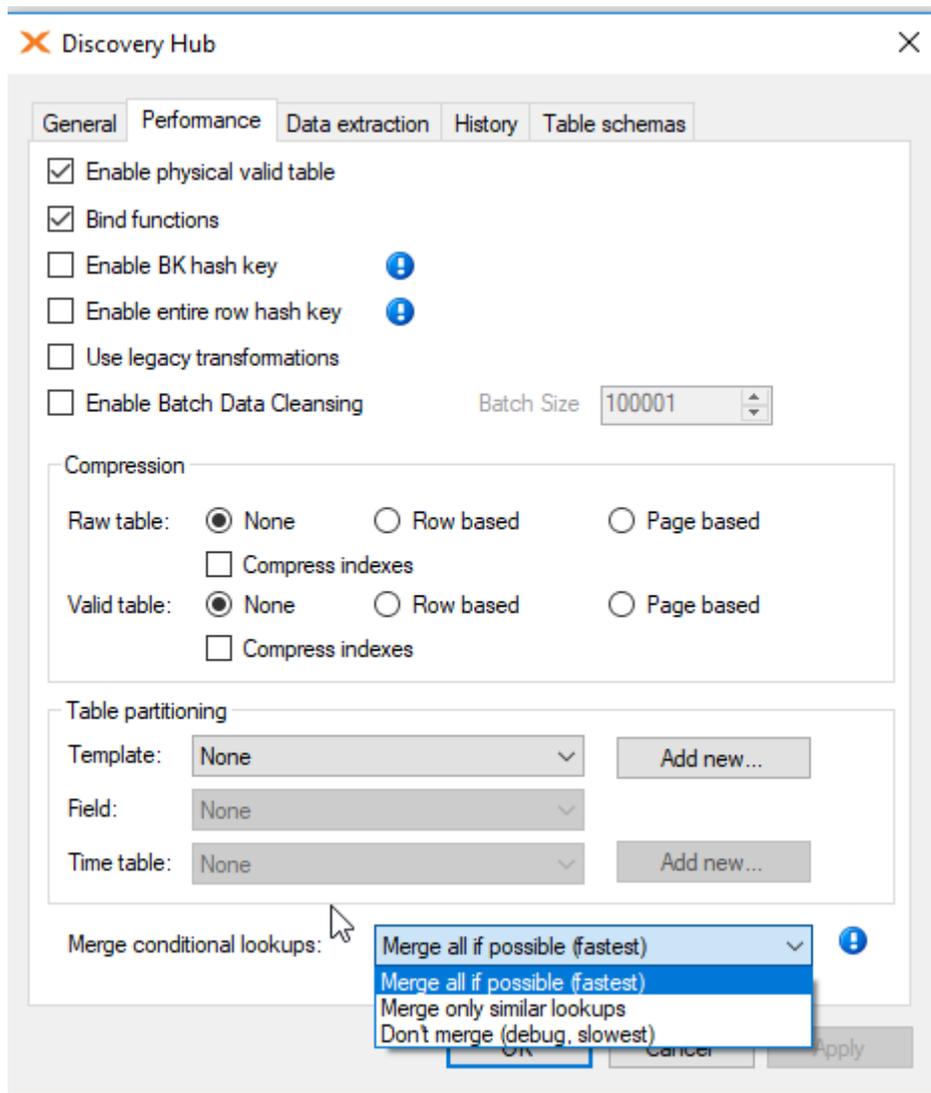




Improved performance: We have optimized the code behind history tables, so all history tables work significantly faster

Improved performance: Possibility to merge all conditional lookups, so all lookups are performed in one go.





Fix: The first relation is automatically created as default, and it is no longer possible to remove default from a relation other than by adding default to another relation or deleting the relation. If there are several relations, one is always default. Thanks to this, adding a lookup without joins can be done by selecting: "Use default relation instead of joins"

Fix: Creating relations by dragging and dropping fields is now possible also if a relation was already created in the opposite table. Relations now work regardless of which table they were created, i.e. they can be used in both tables.

Fix: Consistent dialogs for adding/modifying conditional lookup fields, lookup fields, lookup joins, and lookup sortings

Fix: Allow MAX and MIN with the bit data type (lookups)



Minor General Improvements

Fix: Copying transformations now also copies custom transformations and SQL snippets with parameters

Fix: Data Cleansing of SQL Formatting

Update Connection Manager

The Connection Manager has been modified to also include connections for CData Sources and Semantic Endpoints.

Deprecated Features

Qlik Modeler

The future of modeling – also for Qlik - is the shared semantic layer and we have chosen to deprecate the Qlik Modeler. Users that already have a Qlik Model will experience no differences, but new users will not see the Qlik Tab by default.

