

Qlik Compose May 2022 Release Notes



Skipping versions: Customers who are not upgrading directly from the previous version are strongly encouraged to review the release notes for all versions higher than their currently installed version.

In these release notes:

- *Migration and upgrade (page 2)*
- *What's new? (page 6)*
- *End of Life/Support and Deprecated Features (page 12)*
- *Resolved Issues (page 13)*
- *Known issues (page 35)*

For more information about a particular feature, please refer to the Compose Help.

1 Migration and upgrade

This section describes various upgrade scenarios and considerations.

1.1 Upgrade paths

Compose upgrade path

Direct upgrade is supported from Compose May 2021 or Compose August 2021 only. Customers upgrading from earlier Compose versions need to first upgrade to one of the aforementioned versions and then to Compose May 2022.

Compose for Data Warehouses upgrade path

Compose for Data Warehouses has been superseded by Qlik Compose. Existing Compose for Data Warehouses customers can upgrade to Qlik Compose as described below.



For information on the procedure for upgrading from Compose for Data Warehouses to Compose February 2021, see the [February 2021 release notes](#).

Upgrading from Compose for Data Warehouses 6.6.1 (September 2020) or 7.0 (November 2020):

- a. Upgrade to Compose February 2021.
- b. Upgrade to Compose May 2021.
- c. Compose May 2022.

Upgrading from unsupported Compose for Data Warehouses versions

- Customers upgrading from Compose for Data Warehouses 6.5 or 6.6:
 - a. Upgrade to Compose for Data Warehouses 6.6.1.
 - b. Upgrade to Compose February 2021.
 - c. Upgrade to Compose May 2021.
 - d. Upgrade to Compose May 2022.
- Customers upgrading from Compose for Data Warehouses 6.3 or 6.4:
 - a. Upgrade to Compose for Data Warehouses 6.5.
 - b. Upgrade to Compose for Data Warehouses 6.6.1.
 - c. Upgrade to Compose February 2021.
 - d. Upgrade to Compose May 2021.
 - e. Upgrade to Compose May 2022.
- Customers upgrading from Compose for Data Warehouses 3.1 should contact Qlik Support.

Compose for Data Lakes upgrade path

For information on upgrading from Compose for Data Lakes, see *Migrating from Compose for Data Lakes (page 5)*.

1.2 Required post-upgrade actions after upgrading from Qlik Compose

ETL script enhancements

After upgrading, in order to benefit from the latest enhancements to the task ETL scripts:

- Customers with Data Warehouse projects should regenerate all task ETLs either by selecting the task and clicking the **Generate** button in the **Manage Tasks** and **Manage Data Marts** windows, or by running the generate_project CLI as described in the [Compose online help](#).
- Customers with Data Lake projects should regenerate all task ETLs by selecting the task and clicking the **Generate** button in the **Manage Storage Tasks** window, or by running the generate_project CLI as described in the [Compose online help](#).

Upgrade scripts

After upgrading, depending on the version from which you upgraded, you might need to generate upgrade scripts and run them in your databases.

Upgrade script 1



Should be run only if upgrading from versions earlier than Compose August 2021.

Various performance enhancements require modifications to the internal Compose tables in the following data warehouses:

- Microsoft SQL Server
- Oracle
- Microsoft Azure Synapse Analytics
- Google Cloud BigQuery
- Amazon Redshift

If you have Data Warehouse projects configured to use any of the above databases, you need to generate an upgrade script and then run it in each of the relevant databases.



Running the script in Google Cloud BigQuery and Amazon Redshift databases will delete historical monitoring metadata.

Upgrade script 2



Should be run only if upgrading from versions earlier than Compose August 2021 Service Release 02.

This upgrade script must be run after upgrading, as the database structure has been slightly modified to correctly report the error mart for each source (as part of the *Uniform source consolidation (page 8)* feature).

Upgrade script 3



Should be run only if upgrading from versions earlier than Compose August 2021 SP 12, and only if you have projects with Microsoft Azure Synapse Analytics data warehouse (or intend to create such projects in the future).

Generating and running the upgrade scripts

1. From the **Start** menu, open the **Compose Command Line** console and run the following command:

```
ComposeCli.exe connect
```

2. Run the following command:

```
ComposeCli.exe generate_upgrade_scripts
```

For each of your projects, the CLI output will tell you the name of the script and its location. Each script has a different name, consisting of the script identifier (the bold part), the project name, and a timestamp.

Example of Upgrade script 1:

```
C:\Program Files\Qlik\Compose\data\projects\Project_1\ddl-scripts\ComposeUpgradeFrom2021_5To2021_8Project_1__210714142110.sql
```

Example of Upgrade script 2:

```
C:\Program Files\Qlik\Compose\data\projects\Project_2\ddl-scripts\ComposeUpgradeFrom2021_8SP4To2021_8SP10Project_2__220114142110.sql
```

Example of Upgrade script 3:

```
C:\Program Files\Qlik\Compose\data\projects\Project_3\ddl-scripts\ComposeUpgradeFrom2021_8SP10To2021_8SP12Project_3__220518142110.sql
```

3. Access each of your databases using SQL Workbench or a similar tool and run the script(s).
4. When the script(s) completes successfully, generate and run your tasks in Compose.

1.3 Licensing

Existing Compose for Data Warehouses customers who want to create and manage Data Warehouse projects only in Qlik Compose can use their existing license. Similarly, existing Compose for Data Lakes customers who want to create and manage Data Lake projects only in Qlik Compose can use their existing license.

Customers migrating from Qlik Compose for Data Warehouses or Qlik Compose for Data Lakes, and who want to create and manage *both* Data Warehouse projects and Data Lakes projects in Qlik Compose, will need to obtain a new license. Customers upgrading from Compose February 2021 can continue using their existing license.

It should be noted that the license is enforced only when trying to generate, run, or schedule a task (via the UI or API). Other operations such as Test Connection may also fail if you do not have an appropriate license.

1.4 Migrating from Compose for Data Lakes

Compose for Data Lakes has been superseded by Qlik Compose. Existing Compose for Data Lakes customers can migrate their projects from Qlik Compose for Data Lakes to Qlik Compose. You can migrate both your project definitions and your data although the latter is only required if you need to migrate production data.

For migration instructions, see <https://help.qlik.com/en-US/compose/August2021/pdf/Compose-Release-Notes.pdf>



Migration can be performed from Compose for Data Lakes 6.6 only.

1.5 Compatibility with related Qlik products

Qlik Replicate is required for landing data into the data warehouse or storage while Qlik Enterprise Manager allows you to monitor and control Compose tasks running on different servers. This section lists the supported versions for each of these products.

Compose May 2022 Initial Release

Compose May 2022 Initial Release is compatible with the following Replicate and Enterprise Manager product versions:

- **Qlik Replicate** - Qlik Compose is compatible with Replicate November 2021 latest service release and Replicate May 2022.
- **Enterprise Manager** - Qlik Compose is compatible with Enterprise Manager May 2022.

2 What's new?

The following section describes the enhancements and new features introduced in Qlik Compose May 2022.



The "What's new?" is cumulative, meaning that it also describes features that were already released as part of Compose August 2021 service/patch releases. This is because customers upgrading from initial release versions might not be aware of features that were released in interim service releases.

2.1 What's new in Data Warehouse projects?

The following section describes the enhancements and new features introduced in Qlik Compose Data Warehouse projects.

Keeping changes in the Change Tables

This version introduces a new **Keep in Change Tables** option in the landing zone connection settings:

After applying changes:

Delete from Change Tables

Keep in Change Tables

Archive the Change Tables

Database name: ...

Schema name: ...

When you select the **Keep in Change Tables** option, the changes are kept in the Change Tables after they are applied (instead of being deleted or archived). This is useful as it allows you to:

- Use the changes in multiple Compose projects that share the same landing
- Leverage Change Table data across multiple mappings and/or tasks in the same project
- Preserve the Replicate data for auditing purposes or reprocessing in case of error
- Reduce cloud data warehouse costs by eliminating the need to delete changes after every ETL execution

Referenced dimensions

This version introduces support for referencing dimensions. To facilitate this new functionality, a new **Reference selected dimensions** option has been added to the **Import Dimensions** dialog which, together with the toolbar button, has been renamed to **Import and Reference Dimensions**.

The ability to reference dimensions improves data mart design efficiency and execution flexibility by facilitating the reuse of data sets. Reuse of dimension tables across data marts allows you to break up fact tables into smaller units of work for both design and data loading, while ensuring consistency of data for analytics.

Data mart enhancements

Data mart adjust

This version introduces the following enhancements:

- The automatic data mart adjust feature has been extended to include DROP COLUMN and ADD COLUMN support.
- In previous versions, adding a dimension which did not relate to any fact would require the data mart to be dropped and recreated. From this version, such dimensions can be added using auto-adjust, including Date and Time dimensions.
- The `generate_project` CLI now supports automatic data mart adjust for specific objects. In previous versions, Compose would adjust the data marts by dropping and recreating the tables, regardless of the required change. This would sometimes take a lot of time to complete. From this version, only the changes will be adjusted. For example, if a new column was added to a dimension, only that specific column will be added to the data mart tables. To support this new functionality the `-stopIfDatamartsNeedRecreation` parameter must be included in the command. If this parameter is omitted and the data mart needs to be adjusted, Compose will drop and recreate the data mart tables like it did in previous versions.

Data mart reloading

This version introduces the ability to reload the data mart or parts of the data mart without dropping and recreating it, thereby eliminating costly and lengthy reloading of the data mart while maximizing data availability. Such operations should usually be performed after a column with history has been added by the automatic adjust operation.

To facilitate this, a new `mark_reload_datamart_on_next_run` CLI has been developed. The new CLI allows users to mark dimensions and facts to be reloaded on the next data mart run. These can either be specific dimensions and facts or multiple dimensions and facts (either from the same data mart or different data marts) using a CSV file.

Microsoft Azure Synapse Analytics Enhancements

A number of changes related to statistics have been implemented. In addition, several statements are now tagged with an identifier label for troubleshooting 'problem queries' and identifying possible ways to optimize database settings. Moreover, the addition of labels to ELT queries enables fine-grained workload management and workload isolation via Synapse WORKLOAD GROUPS and CLASSIFIERS.

The identifier labels are as follows:

Table type	Tag
Hubs	CMPS_HubIns

Table type	Tag
Satellites	CMPS_SatIns
Type1 dimensions	CMPS_<data mart name>_DimT1_Init/CMPS_<data mart name>_DimT1_Incr
Type2 dimensions	CMPS_<data mart name>_DimT2_Init/CMPS_<data mart name>_DimT2_Incr
Transactional facts	CMPS_<data mart name>_FctTra_Init/CMPS_<data mart name>_FctTra_Incr
State-oriented facts	CMPS_<data mart name>_FctStO_Init
Aggregated facts:	CMPS_<data mart name>_FctAgg_Init

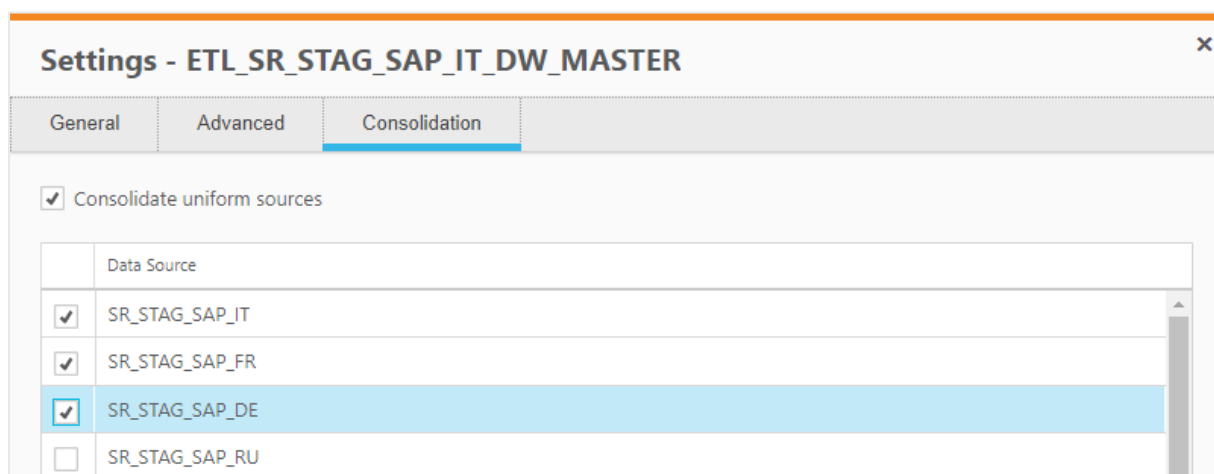
Uniform source consolidation

Uniform source consolidation as its name suggests allows you to ingest data from multiple sources into a single, consolidated, entity.

To enable uniform source consolidation configuration, a new **Consolidation** tab has been added to the data warehouse task settings.

When the **Consolidate uniform sources** option is enabled, Compose will read from the selected data sources and write the data to one consolidated entity. This is especially useful if your source data is managed across several databases with the same structure, as instead of having to define multiple data warehouse tasks (one for each source), you only need to define a single task that consolidates the data from the selected data sources.

Consolidation tab showing selected data sources



Environment variables

Environment variables allow developers to build more portable expressions, custom ETLs, and Compose configurations, which is especially useful when working with several environments such as DTAP (Development, Testing, Acceptance and Production). Different environments (for example, development

and production) often have environment-specific settings such as database names, schema names, and Replicate task names. Variables allow you to easily move projects between different environments without needing to manually configure the settings for each environment. This is especially useful if many settings are different between environments. For each project, you can use the predefined environment variables or create your own environment variables.

Excluding environment variables from export operations

An option has been added to replace environment-specific settings with the defaults when exporting projects (CLI) or creating deployment packages.

To facilitate this functionality, the `--without_environment_specifics` parameter was added to the `export_project_repository` CLI and a **Exclude environment variable values** option was added to the **Create Deployment Package** dialog.

Support for data profiling and data quality rules when using Google Cloud BigQuery

You can now configure data profiling and data quality rules when using Google Cloud BigQuery as a data warehouse.

Attributes case sensitivity support

In previous versions, attempting to create several Attributes with the same name but a different case would result in a duplication error. Now, such attributes will now be created with an integer suffix that increases incrementally for each attribute added with the same name. For example: Sales, SALES_01, and Sales_02.

Associating a Replicate task that writes to a Hadoop target

You can now associate a Replicate task that writes to a Hadoop target with the Compose landing.

Performance improvements

This version provides the following performance improvements:

- Validating a model with self-referencing entities is now significantly faster than in previous versions. For instance, it now takes less than a minute (instead of up to two hours) to validate a model with 5500 entities.
- The time it takes to "Adjust" the data warehouse has been significantly reduced. For instance, it now takes less than three minutes (instead of up to two hours) to adjust a data warehouse with 5500 entities.
- Optimized queries, resulting in significantly improved data warehouse loading and CDC performance.
- Significantly improved the loading speed of data mart Type 2 dimensions with more than two entities. In order to benefit from this improvement, customers upgrading with existing data marts needs to regenerate their data mart ETLs.

Data mart UX improvement

The Data Mart Dimensions tree and the Star Schema Fact tab were redesigned to provide a better user experience.

2.2 What's new in Data Lake projects?

The following section describes the enhancements and new features introduced in Qlik Compose Data Lake projects.

Support for excluding deleted records from ODS views

A **Deleted records in ODS views** section has been added to the **General** tab of the project settings, with the following options:

- **Exclude the corresponding record from the ODS views** - This is the default option as records marked as deleted should not usually be included in ODS views.
- **Include the corresponding record in the ODS views** - Although not common, in some cases, you might want include records marked as deleted in the ODS views in order to analyze the number of deleted records and investigate the reason for their deletion. Also, regulatory compliance might require you to be able to retrieve the past record status (which requires change history as well).



As this was the default behavior in previous versions, you might need to select this option to maintain backward compatibility.

Associating a Replicate task that writes to a Hortonworks Data Platform target

You can now associate a Replicate task that writes to a Hortonworks Data Platform target with the Compose landing connection (in a Cloudera Data Platform (CDP) Compose project).

New Databricks Version


Databricks 9.1 LTS is now supported on all cloud providers (AWS, Azure, and Google Cloud Platform).

2.3 New features common to both Data Warehouse projects and Data Lake projects

New Project title setting

A new **Project title** setting had been added to the **Environment** tab of the project settings. The project title will be shown in the console banner. If both an **Environment Title** and a **Project Title** are defined, the project title will be displayed to the right of the environment title. Unlike the **Environment title** and **Environment type**, which are unique for each environment, the project title is environment independent. This means that the project title will always be retained, even when deploying to a different environment.

The following image shows the banner with both an **Environment title** and a **Project title**:



The banner text is shown without the **Environment title** and **Project title** console labels. This provides greater flexibility as it allows you add any banner text you like, regardless of the actual label name. For example, specifying `Project owner: Mike Smith` in the **Project title** field, will display that text in the banner.

Support for Microsoft Edge Browser

This version introduces support for accessing the Compose console using Microsoft Edge.

Security Hardening

For security reasons, command tasks are now blocked by default. To be able to run command tasks, a Compose administrator needs to turn on this capability using the Compose CLI. For more information, see the Compose online help.



This functionality only applies to command tasks created after a clean installation. If you upgrade to this version, command tasks will continue to work as previously.

3 End of Life/Support and Deprecated Features

This section provides information about End of Life versions, End of Support features, and deprecated features.

- Internet Explorer is no longer supported.

4 Resolved Issues

The following section lists the issues resolved since Compose August 2021 Initial Release.

Jira issue: RECOB-4808

Salesforce case: 2271788

Type: Issue

Component/Process: Environment variables in data mart

Description: After the data mart database name was applied as an environment variable, Compose would not clear the cache automatically, resulting in the old cache object not being reset.

Jira issue: RECOB-4806

Salesforce case: 26263

Type: Issue

Component/Process: UI

Description: Selecting a **Replicate** task would not be possible when using a Hortonworks Data Platform endpoint in a Cloudera Data Platform Compose project.

Jira issue: CMPS-625

Salesforce case: N/A

Type: Enhancement

Component/Process: Environment variables in export

Description: An option has been added to remove environment information when exporting projects (CLI) or creating deployment packages.

To facilitate this functionality, the **--without_environment_specifics** parameter was added to the CLI and a **Replace environment specifics with defaults** option was added to the **Create Deployment Package** window.

Jira issue: RECOB-4822

Salesforce case: 25044

Type: Issue

Component/Process: Project deployment

Description: The following error would sometimes be encountered when deploying a project:

```
Invalid Configuration file the database <name> Landing does not exist
```

Jira issue: RECOB-4802

Salesforce case: 2218782

Type: Enhancement

Component/Process: Project Settings

Description: A new **Project title** field has been added to the project settings **General** tab. The value of the field will be included in the project deployment.

Jira issue: RECOB-4861

Salesforce case: 26682

Type: Issue

Component/Process: Test connection

Description: When the schema name was *, testing the connection for the landing database would return the following error:

```
Object reference not set to an instance of an object
```

Jira issue: RECOB-4854

Salesforce case: 7550

Type: Issue

Component/Process: Lineage

Description: When importing data marts using the **Composecli import_csv** command, the "Show lineage" option for corresponding domain attributes would be disabled.

Jira issue: RECOB-4876

Salesforce case: 27847

Type: Issue

Component/Process: Project Deployment

Description: When a landing connection was removed from the target project, project deployment would fail with the following error:

```
REPO-E-ITMNTFND, Invalid configuration file. The 'Database' 'Landing4' does not exist. REPO,CONFIGURATION_ITEM_NOT_FOUND,Database,Landing4
```

Jira issue: RECOB-4809

Salesforce case: 22405

Type: Issue

Component/Process: Data Marts

Description: Hub tables would sometimes be updated unnecessarily which would result in unnecessary updates of the related dimensions.

Jira issue: RECOB-4836

Salesforce case: N/A

Type: Issue

Component/Process: Data Marts

Description: Failed to set a filter on a dimension or a fact.

Jira issue: RECOB-4779

Salesforce case: 24471

Type: Issue

Component/Process: Data Marts

Description: Filters and expression on dimensions would not work as expected.

Jira issue: RECOB-4882

Salesforce case: 27704

Type: Issue

Component/Process: Data Marts

Description: When a data mart contained an entity with multiple satellites, the query would sometimes be generated incorrectly.

Jira issue: RECOB-4864

Salesforce case: 24810

Type: Issue

Component/Process: Filters and expressions

Description: Tasks with filters or expressions would end with errors.

Jira issue: RECOB-4913

Salesforce case: 27960

Type: Issue

Component/Process: Compare CSV CLI

Description: The Compare CSV CLI would sometimes not complete successfully.

Jira issue: RECOB-4917

Salesforce case: 28209

Type: Issue

Component/Process: Expression Editor

Description: An error would sometimes occur when opening the Expression Editor.

Jira issue: RECOB-4959

Salesforce case: 20574

Type: Issue

Component/Process: Data Warehouse Tasks - Snowflake

Description: Records in the data warehouse would not be updated with a NULL value, even though the data warehouse task was set to "Set the target value to null".

Jira issue: RECOB-4928

Salesforce case: 27075

Type: Issue

Component/Process: Metadata validation in Data Lakes projects

Description: Validating the metadata would fail with an error that "ID" is a reserved word.

Jira issue: RECOB-4722

Salesforce case: 2271788

Type: Issue

Component/Process: Project documentation

Description: In the generated project documentation, the domain name would be shown in the attribute name field.

Jira issue: RECOB-4739

Salesforce case: 22780

Type: Issue

Component/Process: Databricks

Description: After upgrading to 2021.08 SP08, Databricks connection issues would be encountered when a token was revoked.

Jira issue: RECOB-4707

Salesforce case: N/A

Type: Issue

Component/Process: Data Marts - Oracle

Description: The following Oracle syntax error would be encountered during the initial load task command:

::

```
ORA-01400: cannot insert NULL into
```

Jira issue: RECOB-4675

Salesforce case: 15882

Type: Issue

Component/Process: Facts

Description: State oriented facts would not reflect changes that were made to the Type 2 relation or changes that were made to the dimension table.

Jira issue: RECOB-4771

Salesforce case: 24505

Type: Issue

Component/Process: Project deployment

Description: Users with the "Designer" role were not able to deploy project deployment packages.

Jira issue: RECOB-4785

Salesforce case: 10094

Type: Issue

Component/Process: Import CSV

Description: After running the import_csv CLI command to import tasks, the generated task statements would contain a syntax error.

Jira issue: RECOB-4776

Salesforce case: 23553

Type: Issue

Component/Process: Data mart editing

Description: When working with large models, it would not be possible to edit a dimension or fact.

Jira issue: RECOB-4656

Salesforce case: 21696

Type: Issue

Component/Process: CSV Import - Microsoft Azure Synapse Analytics Data Warehouse

Description: Importing a CSV file to a project with a Microsoft Azure Synapse Analytics data warehouse would fail if the CSV contained an NVARCHAR attribute.

Jira issue: RECOB-4666

Salesforce case: 19667

Type: Issue

Component/Process: Security

Description: Resolved security vulnerabilities discovered in Compose 2021.8.0.365.

Jira issue: RECOB-4699

Salesforce case: 23508

Type: Issue

Component/Process: Upgrade Script

Description: Running the generate_upgrade_script command would fail after upgrading to 2021.8.0.425.

Jira issue: RECOB-4045

Salesforce case: 10967

Type: Issue

Component/Process: Generate project CLI

Description: Running the generate_project CLI command with the --database_already_adjusted parameter would drop the Qlik table "TPIL_DMA_RUNNO".

Jira issue: RECOB-3999

Salesforce case: 9804

Type: Issue

Component/Process: Generate project CLI

Description: Running the generate_project CLI command with the --database_already_adjusted parameter would fail with the following error:

SQL compilation error: <p>Object does not exist, or operation cannot be performed.

Jira issue: RECOB-4057

Salesforce case: N/A

Type: Issue

Component/Process: Data Mart

Description: Creating a denormalized new dimension would create the root dimension only.

Jira issue: RECOB-3990

Salesforce case: 2264064

Type: Issue

Component/Process: Workflows

Description: In rare cases, it would not be possible to create, edit, or duplicate workflows.

Jira issue: RECOB-3937, RECOB-3859

Salesforce case: 2236402, 5136

Type: Issue

Component/Process: Upgrade

Description: After migrating to 2021.5, projects containing two domain attributes with the same name but a different case (e.g. abc and Abc) would fail to load with the following error:

`SYS,GENERAL_EXCEPTION, An item with the same key has already been added.`

Jira issue: RECOB-3987

Salesforce case: N/A

Type: Issue

Component/Process: Project Deployment

Description: It would not be possible to open a project after deployment if one schema was missing.

Jira issue: RECOB-4043

Salesforce case: 9043

Type: Issue

Component/Process: Data Mart

Description: Fact tables would contain obsolete VIDs from dimensions, resulting in orphaned records.

Jira issue: RECOB-4033

Salesforce case: 9805

Type: Issue

Component/Process: Data Mart

Description: Data mart loading tasks would sometimes fail with the following error:

Cannot write value for process parameter twice: 1265: Duplicate write to param DimCnt_Tot

Jira issue: RECOB-3204

Salesforce case: 2214622

Type: Issue

Component/Process: Loading data mart dimensions into Snowflake and Microsoft Azure Synapse Analytics

Description: When a data mart ETL task failed, the next task would sometimes load duplicate rows into dimensions.

Jira issue: RECOB-3957

Salesforce case: 2231873

Type: Issue

Component/Process: Data marts

Description: Adding data mart dimensions would sometimes fail without a clear error.

Jira issue: RECOB-3954

Salesforce case: 8634

Type: Issue

Component/Process: Data warehouse validation

Description: The following error would occur when validating the data warehouse:

Index was out of range. Must be non-negative and less than the size of the collection

Jira issue: RECOB-3902

Salesforce case: 7392

Type: Issue

Component/Process: Snowflake

Description: The data warehouse ETL would fail to create a transient table with a "already exists" error.

Jira issue: RECOB-3934

Salesforce case: 8399

Type: Issue

Component/Process: CLI

Description: Importing a project repository to a new project that does not exist it would fail with the following error:

```
Project: 'Project_name' does not exist.
```

Jira issue: RECOB-3636

Salesforce case: 2248515

Type: Issue

Component/Process: Backdating

Description: Backdated data in the Data Warehouse would not get updated in the Data Mart.

Jira issue: RECOB-3703

Salesforce case: 2240557

Type: Issue

Component/Process: Backdating

Description: Migrating a project from an older version would disable the backdating options. The issue was resolved by adding a new CLI command line that sets the "Add actual data row and a precursor row" option for all entities as well as in the project settings.

```
composecli set_backdating_options --project project_name
```

After running the command, refresh the browser to see the changes.

Jira issue: RECOB-3719

Salesforce case: 2260256

Type: Issue

Component/Process: Discovery from Snowflake

Description: When a landing table had a foreign key, discovering the table would result in the following error (excerpt):

Specified argument was out of the range of valid values.

Jira issue: RECOB-3799

Salesforce case: 2264057

Type: Issue

Component/Process: Validation and Schema Evolution

Description: Validation of Databricks storage and Snowflake data warehouse would be excessively long. The slow Databricks validation would also impact schema evolution.

Jira issue: RECOB-4528

Salesforce case: 17678

Type: Issue

Component/Process: Pivot table - Google BigQuery

Description: In Google BigQuery projects, the data mart pivot table displays a "no data error" when there is data in tables.

Jira issue: RECOB-4529

Salesforce case: 17465

Type: Issue

Component/Process: Data profiler - Google BigQuery

Description: In Google BigQuery projects, the following error would be encountered when using the data profiler: "SYS,GENERAL_EXCEPTION,Sequence contains no elements"

Jira issue: RECOB-4535

Salesforce case: 16513

Type: Issue

Component/Process: OID and VID Columns

Description: The OID and VID column names would include the entire path from the fact source to the dimension instead of just the dimension name.

Jira issue: RECOB-4555

Salesforce case: 2260638

Type: Issue

Component/Process: MySQL source

Description: When setting up a MySQL source connection, testing the connection would return the following error: "Object reference not set to an instance of an object".

Jira issue: RECOB-4557

Salesforce case: 19777

Type: Issue

Component/Process: Export CLI

Description: After deleting an entity, export of projects using the CLI would sometimes fail.

Jira issue: RECOB-4584

Salesforce case: 19673

Type: Issue

Component/Process: Data mart loading

Description: When a dimension contained more than 10 entities, loading of the data mart would fail with the following error: "Case expressions may only be nested to level 10.Operation cancelled by user"

Jira issue: RECOB-4595

Salesforce case: 20256

Type: Issue

Component/Process: Data mart task generation

Description: Data mart task generation would fail when attributes of the same entity were assigned to different satellite tables.

Jira issue: RECOB-4633

Salesforce case: 20347

Type: Issue

Component/Process: Bulk Operations

Description: Generating Bulk Operations would not include the last data mart in the list.

Jira issue: RECOB-4636

Salesforce case: 20746

Type: Issue

Component/Process: Data mart loading

Description: Some projects could not be opened after upgrading.

Jira issue: RECOB-4464

Salesforce case: 14522

Type: Issue

Component/Process: CLI

Description: Running the "generate_project" command with the "database_already_adjusted" parameter would reset the data mart to the "Create Tables" state.

Jira issue: RECOB-3917

Salesforce case: 2256585

Type: Issue

Component/Process: Data mart dimensions

Description: Sometimes, rows in dimensions would incorrectly be marked as obsolete.

Jira issue: RECOB-4459

Salesforce case: 17328

Type: Issue

Component/Process: CLI - Export CSV

Description: Running the export_csv command would cause ETL Set generation to fail for lookups with the following error:

SYS,GENERAL_EXCEPTION,startIndex cannot be larger than length of string.<p>Parameter name: startIndex

Jira issue: RECOB-4481

Salesforce case: 17567

Type: Issue

Component/Process: Data marts

Description: Data Mart creation would sometimes fail with the following error "Sequence contains no matching element".

Jira issue: RECOB-4482

Salesforce case: 17567

Type: Issue

Component/Process: Data marts

Description: An error would sometimes be encountered when trying to delete a star schema.

Jira issue: RECOB-4390

Salesforce case: 12810

Type: Issue

Component/Process: ETLs

Description: The ETL for handling data mart dimensions would use the non-optimized approach for one of the statements.

Jira issue: RECOB-4386

Salesforce case: 14640

Type: Issue

Component/Process: Snowflake

Description: After four hours of inactivity, a "Snowflake Authentication token has expired" error would be shown.

Jira issue: RECOB-4500

Salesforce case: 5008

Type: Issue

Component/Process: ETLs

Description: Verification of unused and/or outdated column mapping expressions would lead to redundant errors.

Jira issue: RECOB-4501

Salesforce case: 17659

Type: Issue

Component/Process: Data Marts

Description: Validation of Type 2 dimensions would sometimes fail with an error that no Type 2 columns were detected (and that the dimension should be created as Type 1), even though Type 2 relationships existed in the dimension.

Jira issue: RECOB-4370

Salesforce case: N/A

Type: Issue

Component/Process: Security

Description: Fixes critical vulnerabilities (CVE-2021-45105, CVE-2021-45046, CVE-2021-44228) that may allow an attacker to perform remote code execution by exploiting the insecure JNDI lookups feature exposed by the logging library log4j. The fix replaces the vulnerable log4j library with version 2.16.

Jira issue: RECOB-4293

Salesforce case: 15341

Type: Issue

Component/Process: UI

Description: Editing a data mart entity after creating the data mart would result in all of the fields being reordered alphabetically.

Jira issue: RECOB-4199

Salesforce case: 12178

Type: Issue

Component/Process: Project settings - Snowflake only

Description: Enabling the **Write metadata to the TDWM tables in the data warehouse** option in the project settings would have no effect.

Jira issue: RECOB--4320

Salesforce case: 2160919

Type: Issue

Component/Process: Deployment packages

Description: The source schema connection would not be updated after deploying a deployment package.

Jira issue: RECOB-4258

Salesforce case: 13575

Type: Issue

Component/Process: Data mart

Description: Data mart creation would fail when there were more than 500 relationships.

Jira issue: RECOB-4330

Salesforce case: 13852

Type: Issue

Component/Process: Amazon Redshift

Description: An error would occur when trying to connect to Amazon Redshift using SSL.

Jira issue: RECOB-4351

Salesforce case: 16688

Type: Issue

Component/Process: Data Marts

Description: When there was a 3-tier relationship - for example, Entity_A→Entity_B→Entity_C - and the Fact table contained columns from Entity_A and Entity_C, changes in the relationship values in Entity_B (which should have updated columns from Entity_C in the Fact) would not be updated in the Fact table.

Jira issue: RECOB-4071

Salesforce case: 5258

Type: Issue

Component/Process: Live Views

Description: Reading from live views would take an excessively long time.

Jira issue: RECOB-4387

Salesforce case: 16511

Type: Issue

Component/Process: Microsoft Azure Synapse Analytics

Description: Columns with numeric(n,n) data types would not be retrieved from the Landing Zone.

Jira issue: RECOB-4339

Salesforce case: 5276

Type: Issue

Component/Process: Import

Description: The following error would sometimes be encountered when importing a data mart:

```
SYS,GENERAL_EXCEPTION,Sequence contains no matching element
```

Jira issue: RECOB-4388

Salesforce case: 14522

Type: Issue

Component/Process: ComposeCLI Project Generation

Description: Generating the project would truncate the data mart tables when running the following command:

```
ComposeCli.exe generate_project --project <project name> --database_already_adjusted
```

After generating the project, you need to clear the cache by running the following command:

```
ComposeCli.exe clear_cache --project <project_name> --type storage
```

Jira issue: RECOB-4316

Salesforce case: N/A

Type: Issue

Component/Process: Data Mart Tasks

Description: When loading dimensions, a column would sometimes be used twice, causing the data mart task to fail.

Jira issue: RECOB-4235

Salesforce case: 13170

Type: Issue

Component/Process: Data Mart Tasks

Description: A runtime parameter ("MutCnt_8323" or similar) was incorrectly initialized, causing the data mart task to fail.

Jira issue: RECOB-4104

Salesforce case: 2160919

Type: Enhancement

Component/Process: Microsoft Azure Synapse Analytics Performance

Description: Performance was improved by adding indexes to Transactional and State Oriented fact tables.

Jira issue: RECOB-4105

Salesforce case: 2160919

Type: Enhancement

Component/Process: Microsoft Azure Synapse Analytics Performance

Description: Performance was improved by creating the TEMP table as a HEAP table instead of a HASH table.

Jira issue: RECOB-4106

Salesforce case: 2160919

Type: Enhancement

Component/Process: Microsoft Azure Synapse Analytics Performance

Description: Performance was improved by updating the statistics after each incremental load of the dimensions.

Jira issue: RECOB-4126

Salesforce case: 10967

Type: Enhancement

Component/Process: Microsoft Azure Synapse Analytics Performance

Description: Performance was improved for data mart ETL tasks by adding indexes (over columns used for join clauses) to intermediate tables.

Jira issue: RECOB-4109

Salesforce case: 10247

Type: Issue

Component/Process: Diagnostics

Description: Diagnostic packages would contain the server name of the customer environment, which would sometimes result in users being locked out when the package was deployed in our internal testing environment. Now, the diagnostic packages will be generated without the server name.

Jira issue: RECOB-4113

Salesforce case: 2222648

Type: Issue

Component/Process: Project Documentation

Description: The project documentation for Multi-Table ETLs and Post-Loading ETLs was generated without contents.

Jira issue: RECOB-4142

Salesforce case: 10996

Type: Enhancement

Component/Process: Compose CLI Timeouts

Description: A session expired error would sometimes occur during the CLI commands that took a long time to complete (e.g. `import_csv`). To resolve such timeouts, users can now add the `"-timeout seconds"` parameter to the command. Setting `"--timeout -1"` will run the command without it timing out.

Jira issue: RECOB-3928

Salesforce case: 7892

Type: Issue

Component/Process: Post-ETL Error Reporting

Description: Errors in Post-ETL stored procedures run on Microsoft Azure Synapse Analytics would not be reported.

Jira issue: RECOB-4149

Salesforce case: 2218407

Type: Issue

Component/Process: ETLs on Snowflake

Description: While working with Snowflake via the private link configuration, the engine task would sometimes stop unexpectedly.

Jira issue: RECOB-5239

Salesforce case: 33030

Type: Issue

Component/Process: Data Mart Adjustment

Description: When dropping a relationship to a lookup-table in the Model, adjusting the data mart would fail with the following error:

Object reference not set to an instance of an object

Jira issue: RECOB-5210

Salesforce case: 33745

Type: Issue

Component/Process: Data Mart Task Generation

Description: The following error would sometimes be encountered when generating ETLs after data mart validation:

Sequence contains no matching elements" or "SYS,GENERAL_EXCEPTION,Input string was not in a correct format

Jira issue: RECOB-5217

Salesforce case: 30618

Type: Issue

Component/Process: Data Mart Tasks

Description: Data mart tasks would sometimes fail with the following error:

Invalid object name dbo.TPIL_RUNS

Jira issue: RECOB-4929

Salesforce case: N/A

Type: Enhancement

Component/Process: Data Lakes Project - Real-Time Views

Description: Subquery HIVE errors would sometimes be encountered when creating and reading from the real-time view. The issue was resolved by updating the latest applied partition during runtime,

Jira issue: RECOB-5081

Salesforce case: 26461

Type: Issue

Component/Process: Satellite Loading Performance

Description: Performance issues would sometimes be encountered when loading data warehouse satellites tables.

Jira issue: RECOB-5064

Salesforce case: 29989

Type: Issue

Component/Process: Project documentation

Description: When generating project documentation, the following error would sometimes occur:

System.OutOfMemoryException

Jira issue: RECOB-5137

Salesforce case: 30948

Type: Issue

Component/Process: Adding dimensions

Description: Adding a dimension without the "dummy" row would result in incomplete loading on the next task run.

5 Known issues

This section describes the known issues for this release.

Jira issue: N/A

Salesforce case: N/A

Component/Process: Schema Evolution - New Columns

Description: When using Replicate to move source data to Compose, both the Full Load and Store Changes replication options must be enabled. This means that when Replicate captures a new column, it is added to the Replicate Change Table only. In other words, the column is stored without being added to the actual target table (which in terms of Compose is the table containing the Full Load data only i.e. the landing table).

For example, let's assume the **Employees** source table contains the columns **First Name** and **Last Name**. Later, the column **Middle Name** is added to the source table as well. The Change Table will contain the new column while the Replicate Full Load target table (the Compose Landing table) will not.

In older versions of Compose for Data Warehouses, mappings relied on the Full Load tables (the Compose Landing tables), meaning that users were not able to see any new columns (i.e. **Middle Name** in the above example) until they were created in the Full Load tables via a reload.

From Compose May 2021, the Compose Discover and Mappings windows show changes to new columns that exist in both the Change Tables and the Replicate Full Load target tables. This allows Schema Evolution to suggest adding columns that exist in either of them.

Although this is a much better implementation, it may create another issue. If a Full Load or Reload occurs in Compose before the Replicate reload, Compose will try to read from columns that have not yet been propagated to the Landing tables (assuming they exist in the Change Tables only). In this case, the Compose task will fail with an error indicating that the columns are missing.

Should you encounter such a scenario, either execute a reload in Replicate or create an additional mapping without the new columns to allow Compose to perform a Full Load from the Landing tables.

Jira issue: N/A

Salesforce case: N/A

Component/Process: Referenced dimensions

Description: If a dimension being referenced is dropped and created, or reloaded for any reason (for example, the source data mart is fully rebuilt on each load), any facts to which the referenced dimension was added should be reloaded too. Currently, Compose does not handle this automatically.

Workaround:

Run the data marts containing the referenced dimensions.

Jira issue: RECOB-5315

Salesforce case: 33522

Component/Process: Snowflake Data Warehouse Tasks

Description: When generating the data warehouse task, if any attribute with the JSON data type is defined as Type 2, the following error will occur:

```
SYS,GENERAL_EXCEPTION,invalid enum value<p>Parameter name: ACDataType
```