



Informatica® Cloud Data Integration
April 2024

What's New

Informatica Cloud Data Integration What's New
April 2024

© Copyright Informatica LLC 2016, 2024

This software and documentation are provided only under a separate license agreement containing restrictions on use and disclosure. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC.

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation is subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License.

Informatica, Informatica Cloud, Informatica Intelligent Cloud Services, PowerCenter, PowerExchange, and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and many jurisdictions throughout the world. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners.

Portions of this software and/or documentation are subject to copyright held by third parties. Required third party notices are included with the product.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, report them to us at infa_documentation@informatica.com.

Informatica products are warranted according to the terms and conditions of the agreements under which they are provided. INFORMATICA PROVIDES THE INFORMATION IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

Publication Date: 2024-04-05

Table of Contents

Preface	5
Informatica Resources.	5
Informatica Documentation.	5
Informatica Intelligent Cloud Services web site.	5
Informatica Intelligent Cloud Services Communities.	5
Informatica Intelligent Cloud Services Marketplace.	5
Data Integration connector documentation.	6
Informatica Knowledge Base.	6
Informatica Intelligent Cloud Services Trust Center.	6
Informatica Global Customer Support.	6
Chapter 1: April 2024	7
Important notices.	7
New features and enhancements.	8
Changed behavior	10
Connectors.	12
Enhanced connectors.	12
Changed behavior	17
Chapter 2: February 2024	20
New features and enhancements.	20
Changed behavior	21
Connectors	22
Enhanced connectors.	22
Changed behavior.	22
Chapter 3: December 2023	24
Connectors	24
Enhanced connectors.	24
Chapter 4: November 2023	25
Important notices	25
New features and enhancements	26
Changed behavior.	27
Connectors.	28
Enhanced connectors.	28
Changed behavior.	30
Chapter 5: Upgrade	32
Preparing for the upgrade.	32

Disk space requirements and file preparation.	32
Dynamic files in preprocessing commands.	33
Post-upgrade tasks for the February 2024 release.	33
Hive Connector.	33
Index.	34

Preface

Read *What's New* to learn about new features, enhancements, and behavior changes in Informatica Intelligent Cloud Services™ Data Integration for the April 2024 release. You can also learn about upgrade steps that you might need to perform.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Documentation

Use the Informatica Documentation Portal to explore an extensive library of documentation for current and recent product releases. To explore the Documentation Portal, visit <https://docs.informatica.com>.

If you have questions, comments, or ideas about the product documentation, contact the Informatica Documentation team at infa_documentation@informatica.com.

Informatica Intelligent Cloud Services web site

You can access the Informatica Intelligent Cloud Services web site at <http://www.informatica.com/cloud>. This site contains information about Informatica Cloud integration services.

Informatica Intelligent Cloud Services Communities

Use the Informatica Intelligent Cloud Services Community to discuss and resolve technical issues. You can also find technical tips, documentation updates, and answers to frequently asked questions.

Access the Informatica Intelligent Cloud Services Community at:

<https://network.informatica.com/community/informatica-network/products/cloud-integration>

Developers can learn more and share tips at the Cloud Developer community:

<https://network.informatica.com/community/informatica-network/products/cloud-integration/cloud-developers>

Informatica Intelligent Cloud Services Marketplace

Visit the Informatica Marketplace to try and buy Data Integration Connectors, templates, and mapplets:

<https://marketplace.informatica.com/>

Data Integration connector documentation

You can access documentation for Data Integration Connectors at the Documentation Portal. To explore the Documentation Portal, visit <https://docs.informatica.com>.

Informatica Knowledge Base

Use the Informatica Knowledge Base to find product resources such as how-to articles, best practices, video tutorials, and answers to frequently asked questions.

To search the Knowledge Base, visit <https://search.informatica.com>. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Intelligent Cloud Services Trust Center

The Informatica Intelligent Cloud Services Trust Center provides information about Informatica security policies and real-time system availability.

You can access the trust center at <https://www.informatica.com/trust-center.html>.

Subscribe to the Informatica Intelligent Cloud Services Trust Center to receive upgrade, maintenance, and incident notifications. The [Informatica Intelligent Cloud Services Status](#) page displays the production status of all the Informatica cloud products. All maintenance updates are posted to this page, and during an outage, it will have the most current information. To ensure you are notified of updates and outages, you can subscribe to receive updates for a single component or all Informatica Intelligent Cloud Services components. Subscribing to all components is the best way to be certain you never miss an update.

To subscribe, on the [Informatica Intelligent Cloud Services Status](#) page, click **SUBSCRIBE TO UPDATES**. You can choose to receive notifications sent as emails, SMS text messages, webhooks, RSS feeds, or any combination of the four.

Informatica Global Customer Support

You can contact a Global Support Center through the Informatica Network or by telephone.

To find online support resources on the Informatica Network, click **Contact Support** in the Informatica Intelligent Cloud Services Help menu to go to the **Cloud Support** page. The **Cloud Support** page includes system status information and community discussions. Log in to Informatica Network and click **Need Help** to find additional resources and to contact Informatica Global Customer Support through email.

The telephone numbers for Informatica Global Customer Support are available from the Informatica web site at <https://www.informatica.com/services-and-training/support-services/contact-us.html>.

CHAPTER 1

April 2024

The following topics provide information about new features, enhancements, and behavior changes in the April 2024 release of Informatica Intelligent Cloud ServicesSM Data Integration.

Important notices

Informatica has modified the levels of customer support that it offers for one or more Informatica Intelligent Cloud Services Data Integration features in the current release. The following notices identify the features and the current support levels.

Preview lifted

Effective in the April 2024 release, the following connectors are lifted from preview and are production ready:

- Microsoft Fabric Data Warehouse
- Microsoft Fabric Lakehouse
- Microsoft Fabric OneLake

Effective in the April 2024 release, the following functionality is lifted from preview and is production ready:

Intelligent structure models

You can use Cobol copybooks input type to create models in Intelligent Structure Discovery.

For more information, see *Intelligent structure models* .

Mappings in SQL ELT mode

Create a mapping in SQL ELT mode when your source and target are in the same cloud ecosystem and you want all the mapping logic to be processed by your cloud ecosystem. For example, you need to move data in an Azure Data Lake Storage data lake to your Snowflake cloud data warehouse, and you want all the data processing to occur within Snowflake.

For more information about mappings in SQL ELT mode, see *Mappings*.

Deprecation initiated

Effective in the April 2024 release, the following functionality enters deprecation:

- Use a Visio template to define parameterized data flow logic that you can use in mapping tasks.

While you can use the functionality in the current release, Informatica intends to drop support for the functionality in the July 2024 release.

New features and enhancements

Informatica Intelligent Cloud Services Data Integration includes the following new features and enhancements in this release.

New UI experience

The Data Integration user interface has a new look and feel that presents the features you'll use in greater clarity.

New mapping task editor

The updated user interface includes a new mapping task editor. The new editor provides a more intuitive design flow and allows you to edit multiple assets at the same time.

The properties that you configure in the mapping task are the same in the new editor. However the page that you use to configure the properties might be different. For example, you now configure lookup connection and object parameters on the **Sources** page.

Some pages of the new task editor have new names. For example, the **Schedule** page is now the **Runtime Options** page.

For more information about mapping tasks, see *Tasks*.

Asset tags

When you export assets, asset tags are included in the export package and imported with assets to the target project. You can choose to merge tags with tags in the target project when you overwrite existing assets.

You can also merge asset tags when you copy an asset to a new location and replace an existing asset with the same name.

For more information about asset tags, see *Asset Management*.

Configure options for multiple flat file delimiters

When you use a flat file in a Source, Target, or Lookup transformation and you specify multiple delimiter characters, you can choose to treat the characters as a single delimiter or multiple delimiters. By default, Data Integration treats the characters as a single delimiter.

For more information, see *Transformations*.

Data transfer tasks

When you use a database connection, you can define a custom query as the first source in a data transfer task.

For more information about data transfer tasks, see *Tasks*.

FHIR validation in the Data Services transformation

The Data Services transformation can validate incoming FHIR resources in JSON format.

For more information, see *Transformations*.

Hierarchical mapper

You can create a hierarchical mapper to map a hierarchical schema, such as the schema for an industry-standard message, to another hierarchical schema.

For more information, see *Components*.

Mappings in SQL ELT mode

Mappings in SQL ELT mode have the following enhancements:

- When you create a mapping in SQL ELT mode, you can preview data for individual transformations to test the mapping logic. Preview the data for a transformation on the **Preview** panel of the transformation. Select the number of source rows to process and the runtime environment that runs the preview job.
- When you create a mapping in SQL ELT mode, you can preview the SQL query that Data Integration pushes to the cloud data warehouse. Preview the query in the **SQL ELT Query** panel in the Mapping Designer.

For more information about mappings in SQL ELT mode, see *Mappings*.

New B2B Transformation step in mappings

You can use the new B2B transformation step in a mapping to update an event status. You can select a dynamic event status name that is configured in the field mapping or select an event status name that is configured in B2B Gateway.

For more information about the B2B transformation step, see *Mappings*.

Platform REST API

This release includes the following enhancements to the REST API:

Source control

You can get repository connection details such as the repository connection ID for a given project ID or project name using the repositoryConnection resource.

Additional object specifications in import API resource

In an import POST request, you can use key-value pairs in the additionalProviderFlags field to include additional object specifications such as in-out parameter values and sequence generator values.

For more information, see *REST API Reference*.

Serverless runtime environment

You can use serverless runtime environments configured in Azure Native ISV Services.

For a list of connectors that can use Azure serverless runtime environments, see *Runtime environments*.

Source control

This release includes the following enhancements to source control:

Commit tags

When you check in assets to your source control repository, you can add a commit tag to assets in the check-in.

Cloud-hosted Bitbucket repositories

You can use a cloud-hosted Atlassian Bitbucket repository for source-controlled Data Integration assets.

Project-level Git repositories

You can configure a Git repository for each project instead of using the global Git repository that is configured for the organization. To use this feature, the organization administrator must enable it.

For more information, see *Asset Management*.

User-defined function parameters

You can parameterize a user-defined function or its arguments with an in-out parameter that's configured as an expression variable.

For more information about user-defined functions, see *Components*. For more information about in-out parameters, see *Mappings*.

Variables in advanced mode

You can use data filter variables and system variables in mappings in advanced mode.

For more information about data filter variables, see *Tasks*. For more information about system variables, see *Function Reference*.

Changed behavior

Informatica Intelligent Cloud Services Data Integration includes the following changed behaviors in this release.

Job details for mappings in SQL ELT mode

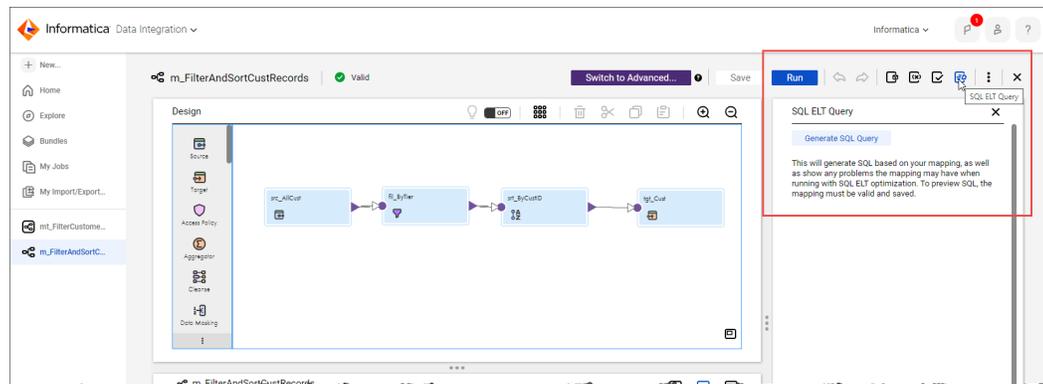
When you run a mapping in SQL ELT mode, the job details page shows the individual target results for each target in the job.

Previously, the job details page showed only one target and the number of target rows shown was cumulative for all targets in the mapping.

For more information about viewing job results, see *Monitor*.

SQL ELT query preview

When you preview the SQL query for a mapping that uses SQL ELT optimization, the query panel is renamed to **SQL ELT Query**, and the button that you click to generate the query is labeled **Generate SQL Query** as shown in the following image:



Previously, the panel was named **SQL ELT Optimization**, and the button was labeled **SQL ELT Preview**.

For more information about previewing SQL ELT queries for mappings, see *Mappings*.

Rounding for double data type values

For double data type source fields that are connected to a downstream decimal or string data type field, Data Integration retains digits beyond the 15th digit.

Previously, Data Integration rounded all double data type values to the 15th digit.

For more information about double data types, see *Function Reference*.

Taskflows

Taskflows include the following changed behaviors:

- In a Data Task step that you use to add a mapping task, the **Path** column corresponding to the in-out parameters in the input and output fields includes the mapping task ID. Previously, the **Path** column included only the mapping task name and not the mapping task ID.

Note: A taskflow that uses mapping tasks with the same name but different in-out parameters becomes invalid. You must re-select the mapping tasks in the taskflow to fix the issue. However, the taskflow will still run successfully even if it shows up as invalid.

- In a subtaskflow, when a Data Task step that uses a mapping task fails, and you select the Subtaskflow step in the taskflow instance from the **All Jobs** or **My Jobs** page, the **Source** section on the **Fault** tab displays the subtaskflow name followed by the mapping task name that failed. If a subtaskflow containing multiple Data Task steps fails, the **Source** section on the **Fault** tab displays the subtaskflow name followed by the mapping task name of the first Data Task step that failed.

With this change, you can easily identify issues with the subtaskflow and mapping task.

The following image shows the subtaskflow name and the mapping name in the **Source** section on the **Fault** tab:

The screenshot shows a taskflow instance titled 'PTaskflow10-949588242813112320' which has failed. The design view shows a subtaskflow 'TF_With_MCT' (indicated by a red box and a red 'X') that has failed. The 'Fault' tab is selected, showing the 'Source' section with the value 'TF_With_MCT : MCT_Basic'.

Step	Duration (HH:MM:SS.sss)
Start	00:00:00.009
TF_With_MCT	00:00:06.350
Notification Task 1	00:00:00.416
End	00:00:00.003

Name	Value
Code	faultDueToNotRunOrFailed
Reason	There is no active Secure Agent in the Secure Agent group invr84ml02.informatica.com or the 'Data Integration Server' service is user stopped or disabled. Start at least one Secure Agent belonging to the Secure Agent group or enable the 'Data Integration Server' service.
Source	TF_With_MCT : MCT_Basic

Previously, in a subtaskflow, when a Data Task step with a mapping task failed, the XPath for the subtaskflow was displayed in the **Source** section on the **Fault** tab, which was not useful.

The following image shows the XPath of the subtaskflow in the **Source** section on the **Fault** tab:

The screenshot shows a taskflow instance titled 'PTaskflow-949594968163729408' which has failed. The design view shows a subtaskflow 'TF_With_MCT' (indicated by a red box and a red 'X') that has failed. The 'Fault' tab is selected, showing the 'Source' section with the XPath value: '/process/Row/scope[@name='the5egcl']/scope[@name='TF_With_MCT']/sequence/invoker[@name='TF_With_MCT']'.

Step	Duration (HH:MM:SS.s)
Start	00:00:00.0...
TF_With_MCT	00:00:05.4...
Notification Task 1	00:00:00.0...
End	00:00:00.0...

Name	Value
Code	faultDueToNotRunOrFailed
Reason	Unexpected exception found.
Source	/process/Row/scope[@name='the5egcl']/scope[@name='TF_With_MCT']/sequence/invoker[@name='TF_With_MCT']

For more information, see *Taskflows*.

Connectors

The April 2024 release includes the following enhanced connectors.

Enhanced connectors

This release includes enhancements to the following connectors.

Adabas Connector

- You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.
- You can override source objects through parameterization. Parameterization provides the ability to parameterize the number of objects and change the object names.

Adabas CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Amazon Athena Connector

You can configure AWS Identity and Access Management (IAM) authentication when the Secure Agent is installed on an Amazon Elastic Compute Cloud (EC2) system.

Amazon Redshift V2 Connector

This release includes the following enhancements for Amazon Redshift V2 Connector:

- **SQL ELT optimization enhancement**
 - When you run a mapping enabled for SQL ELT optimization and the mapping contains multiple pipelines, you can define the flow run order when you want Data Integration to load the targets from the pipelines in a particular order.
 - When you configure a mapping to write to multiple targets in a single pipeline, with each target configured for any write operation, the target operations are processed in a predefined order of delete, update, upsert, and insert.
- You can configure a SQL transformation to process a user-entered query to call a stored procedure that returns a value or no values.
- You can configure a virtual private cloud endpoint to enable private communication between the Secure Agent and the Amazon S3 bucket for staging.

Databricks Delta Connector

This release includes the following enhancements for Databricks Delta Connector:

- **SQL ELT optimization enhancement**

When you run a mapping enabled for SQL ELT optimization and the mapping contains multiple pipelines, you can define the flow run order when you want Data Integration to load the targets from the pipelines in a particular order.
- You can configure partitioning to optimize the mapping performance at run time when you write data to Databricks Delta targets.

- You can improve the staging performance of a mapping that writes to Databricks Delta.
- You can configure a connection and mapping in one environment and then migrate and run the mapping in another environment.

Db2 Data Map Connector

- You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.
- You can override source objects through parameterization. Parameterization provides the ability to parameterize the number of objects and change the object names.

Db2 for i CDC Connector

This release includes the following enhancements for Db2 for i CDC Connector:

- You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.
- The optional mapping source property, **Journal Name**, enables you to override the library and journal name in the PowerExchange CAPI_CONNECTION statement.

Db2 for LUW CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Db2 for z/OS Bulk Load Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Db2 for z/OS CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Db2 for z/OS Image Copy Connector

- You can override source objects through parameterization. Parameterization provides the ability to parameterize the number of objects and change the object names.
- You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Db2 for z/OS Unload File Connector

- You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.
- You can override source objects through parameterization. Parameterization provides the ability to parameterize the number of objects and change the object names.

DB2 Loader Connector

You can encrypt passwords when connecting to DB2 databases.

Google BigQuery V2 Connector

This release includes the following enhancements for Google BigQuery V2 Connector:

- You can create a target table with clustered columns at run time.
- You can configure the threshold number of errors in a mapping task. The mapping continues to run until the number of errors reaches the set threshold value, then it fails.

- You can enable sorted input in a cached Lookup transformation to enhance the lookup performance.
- You can add a label in a mapping to filter the associated jobs in the Google Log Explorer.
- You can configure views in a mapping in the following scenarios:
 - Within a Lookup transformation.
 - Within a Source transformation with staging optimization enabled.

IMS Connector

This release includes the following enhancements for IMS Connector:

- You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.
- You can override source objects through parameterization. Parameterization provides the ability to parameterize the number of objects and change the object names.

IMS CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

JDBC V2 Connector

You can use a JDBC V2 connection to read data from Salesforce Data Cloud.

Kafka Connector

You can configure the polling interval value in the source advanced properties. The polling interval determines how long the agent waits to retrieve records from the Kafka broker.

Microsoft Azure Cosmos DB SQL API Connector

When you write data in bulk to Microsoft Azure Cosmos DB SQL API, you can set the mapping to throttle throughput of low priority requests and prioritize higher priority requests.

Microsoft Azure Synapse SQL Connector

This release includes the following enhancements for Microsoft Azure Synapse SQL Connector:

- **SQL ELT optimization enhancement**
When you run a mapping enabled for SQL ELT optimization and the mapping contains multiple pipelines, you can define the flow run order when you want Data Integration to load the targets from the pipelines in a particular order.
- When you run a mapping that reads from and writes to Microsoft Azure Synapse SQL, you can override the schema name that the Secure Agent uses to create external tables for staging data files.
- You can use the ISNULL or ISNOTNULL operators when you define a source filter or an uncached lookup condition.
- You can push the IS_DATE() function to process in Microsoft Azure Synapse SQL.

Microsoft Fabric Data Warehouse Connector

This release includes the following enhancements for Microsoft Fabric Data Warehouse Connector:

- You can use custom queries to read data from Microsoft Fabric Data Warehouse source objects.
- You can configure a connection and mapping in one environment and then migrate and run the mapping in another environment.
- You can determine the behavior of a task that encounters errors while executing pre-SQL or post-SQL commands.

- You can perform delete and data driven operations when you write to a Microsoft Fabric Data Warehouse target.
- You can truncate the target data before inserting the data to a Microsoft Fabric Data Warehouse target.

Microsoft Fabric Lakehouse Connector

This release includes the following enhancements for Microsoft Fabric Lakehouse Connector:

- You can configure a connection and mapping in one environment and then migrate and run the mapping in another environment.
- You can use custom queries to read data from Microsoft Fabric Lakehouse source objects.

Microsoft Fabric OneLake Connector

This release includes the following enhancements for Microsoft Fabric OneLake Connector:

- You can read files from directories and sub-directories.
- You can use wildcard characters to specify the source files and source directory names.
- You can configure an authenticated or unauthenticated proxy server to connect to Microsoft Fabric OneLake.
- You can read data from semi-structured or unstructured source files in a mapping in advanced mode.
- You can append data to an existing file in Microsoft Fabric OneLake.

Microsoft SQL Server Connector

This release includes the following enhancements for Microsoft SQL Server Connector:

- You can run Microsoft SQL Server mappings using the upgraded JDBC driver version 6.0.0.001282 or ODBC driver version 08.02.1500.
- You can configure the connection timeout value for test connection while creating a Microsoft SQL Server connection.

Microsoft SQL Server CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

MySQL Connector

You can configure the connection timeout value for test connection while creating a MySQL connection.

MySQL CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

ODBC Connector

This release includes the following enhancements for ODBC Connector:

- You can read from or write data to Sybase ASE databases.
- You can use Kerberos authentication to connect to SSL-enabled Sybase ASE databases.

Oracle Cloud Object Storage Connector

You can read and write Gzip-compressed flat files.

Oracle CDC V2 Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Oracle Connector

This release includes the following enhancements for Oracle Connector:

- You can run Oracle mappings using the upgraded JDBC driver version 6.0.0.001165 or ODBC driver version 08.02.3072.
- You can configure the connection timeout value for test connection while creating an Oracle connection.

PostgreSQL Connector

You can use Kerberos authentication to connect to PostgreSQL databases.

PostgreSQL CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Salesforce Connector

You can use version 59.0 and 60.0 of Salesforce API to create a Salesforce connection and access Salesforce objects.

SAP HANA Connector

You can read from SAP HANA external views that are created by SAP BW CompositeProviders in the SAP system.

SAP HANA CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Sequential File Connector

This release includes the following enhancements for Sequential File Connector:

- You can set a user-defined parameter in the Custom Properties connection property that allows you to specify connection property overrides through a parameter file.
- You can override source objects through parameterization. Parameterization provides the ability to parameterize the number of objects and change the object names.

Snowflake Data Cloud Connector

This release includes the following enhancements for Snowflake Data Cloud Connector:

- **Mapping in SQL ELT mode enhancements**
 - You can preview data for individual transformations to validate the mapping logic. You can also preview data from fields that contain primitive and hierarchical data types.
 - You can preview the SQL queries at design time.
 - When you configure multiple targets, you can view statistics for each target separately in the job details page.
 - You can push a few data generation, encryption, string and binary, and numeric functions to Snowflake. For the list of supported functions, see the Snowflake Data Cloud Connector documentation.
- When you run a mapping enabled for SQL ELT optimization and the mapping contains multiple pipelines, you can define the flow run order when you want Data Integration to load the targets from the pipelines in a particular order.
- You can set the default column value for the return field in a Lookup transformation in a Snowflake mapping. This feature does not apply to mappings in advanced mode.

- You can override the `query_tag` parameter that is defined in the Snowflake connection in a mapping task. This feature does not apply to mappings in advanced mode.
- You can remove the leading and trailing white spaces from strings when you write to a target. You can also convert fields that contain only spaces into empty strings and write them as Null values to a target. This feature does not apply to mappings in advanced mode.
- You can read from Apache Iceberg tables that are managed by Snowflake or any external catalog, and write to Apache Iceberg tables that are managed by Snowflake.

Preview Notice: Reading from Apache Iceberg tables that are managed by Snowflake or any external catalog, and writing to Apache Iceberg tables that are managed by Snowflake is available for preview. Preview functionality is supported for evaluation purposes but is unwarranted and is not supported in production environments or any environment that you plan to push to production. Informatica intends to include the preview functionality in an upcoming release for production use, but might choose not to in accordance with changing market or technical circumstances. Note that if you are working on a preview POD, all data is excluded from SOC 2 compliance coverage. For more information, contact Informatica Global Customer Support.

Tableau V3 Connector

You can configure a personal access token authentication to connect to Tableau Server or Tableau Online.

VSAM Connector

This release includes the following enhancements for VSAM Connector:

- You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.
- You can override source objects through parameterization. Parameterization provides the ability to parameterize the number of objects and change the object names.

VSAM CDC Connector

You can set a user-defined parameter in the **Custom Properties** connection property that allows you to specify connection property overrides through a parameter file.

Changed behavior

This release includes changes in behavior for the following connectors.

Amazon Redshift V2 Connector

When you configure a mapping enabled for full SQL ELT optimization and define a simple filter on a column of Super data type, the mapping fails.

Previously, the mapping ran with full SQL ELT optimization. However, zero records were written to target.

Google BigQuery V2 Connector

Effective in this release, Google BigQuery V2 Connector includes the following changes in behavior:

- In an existing mapping task configured with the **Stop on errors** advanced session property, the value is honored. The mapping continues to run until the number of errors reaches the set threshold value, then it fails. Previously, the mapping failed at the first occurrence of an error, regardless of the value set for the **Stop on errors** advanced session property.

- If you enabled `-DENABLE_NULL_FLAG_FOR_UNCACHED_LOOKUP=true` property in the JVM options of the Secure Agent and have an existing mapping with the following scenarios, the mapping fetches correct data in the first row of the table:

- Configured an uncached Lookup transformation to read from Google BigQuery.
- Lookup transformation input contains null data in the first row.

Previously, the mapping fetched incorrect data in the first row of the table.

- If you enabled `-DENABLE_SORTED_INPUT_FOR_LKP=true` property in the JVM options of the Secure Agent and have an existing mapping with the following scenarios, the mapping fails:

- Staging is not selected as the read mode.
- The **Sorted Input** property is selected and the **Use EXPORT DATA Statement to stage** property is not selected in the lookup advanced properties.

Previously, the mapping ran successfully.

- When a Google BigQuery table is imported as a single object source in Staging mode using CSV format, and the backend schema in Google BigQuery is modified by rearranging the columns, the mapping correctly writes the data to the corresponding columns in the target table.

Previously, the mapping either failed to run or corrupted the data in the target table.

Microsoft Azure Synapse SQL Connector

Effective in this release, Microsoft Azure Synapse SQL Connector includes the following change in behavior:

- When you choose a storage type in the connection properties, the mapping shows the properties relevant to that storage type in both the source and target properties, simplifying the user experience. Previously, all storage properties were displayed.
- The mapping displays the advanced runtime properties only if the configured runtime environment is up and running. If you have configured a Secure Agent group, then the mapping displays the advanced runtime properties if at least one agent in the Secure Agent group is up and running. Previously, the mapping displayed the advanced runtime properties irrespective of the runtime environment status.
- If you configure an uncached Lookup transformation to look up data from a Microsoft Azure Synapse SQL source, where the source data of the string data type contains empty values, the uncached lookup query condition uses the correct values from the lookup object. Previously, the uncached lookup query condition used incorrect values from the lookup object.

Microsoft SQL Server Connector

Effective in this release, if you use the latest DataDirect driver and use SSL encryption, the OpenSSL version 3.0 is used by default.

Previously, the default version was OpenSSL 1.1.1.

If you want to use an earlier version of OpenSSL, add the following property in the **Runtime Advanced Connection Properties** field:

```
AllowedOpenSSLVersions=<OpenSSL_version>
```

Oracle Connector

Effective in this release, if you use the latest DataDirect driver and use SSL encryption, the OpenSSL version 3.0 is used by default.

Previously, the default version was OpenSSL 1.1.1.

If you want to use an earlier version of OpenSSL, add the following property in the **Runtime Advanced Connection Properties** field:

```
AllowedOpenSSLVersions=<OpenSSL_version>
```

Snowflake Data Cloud Connector

Effective in this release, when you configure a mapping enabled for full SQL ELT optimization and specify a temporary key column that contains the Binary data type in the **Update Column** field for the target, the mapping runs successfully with full SQL ELT optimization.

Previously, the mapping ran without SQL ELT optimization.

CHAPTER 2

February 2024

The following topics provide information about new features, enhancements, and behavior changes in the February 2024 release of Informatica Intelligent Cloud ServicesSM Data Integration.

New features and enhancements

Informatica Intelligent Cloud Services Data Integration includes the following new features and enhancements in this release.

You can also watch the following video to learn about some of the new features and enhancements in the February 2024 release:



Data Integration REST API

You can include the federated ID in an fwConfig GET request so that the response includes the project and folder location of the fixed-width file format. Use the following URI: `/api/v2/fwConfig/frs/<id>`

For more information, see *REST API Reference*.

Hierarchical schemas based on ZIP files

You can upload a ZIP file that contains one or more XSD schema files to create a hierarchical schema.

For more information, see *Components*.

Importing assets with persisted values

When you import assets that contain in-out parameters and sequence generators, you can choose to either reset, retain, or overwrite the persisted value.

For more information about importing assets, see *Asset Management*.

Output group settings for intelligent structure models

When you create an intelligent structure model based on an XML or XSD file, you can configure how Intelligent Structure Discovery defines the output groups before you discover the structure of the file. You can choose from the following options:

- Default output groups. Can generate normalized or denormalized output groups. This method was used for all intelligent structure models prior to this release.
- Normalized hierarchy relationships. Generates output groups for each repeating element. This method is generally faster and uses less memory than generating default output groups.
- Entity relationships. Generates output groups for each repeating element and for complex data types. This method creates the most compact representation of complex data types.

These new options replicate the options available in PowerCenter to help you migrate assets from PowerCenter to Informatica Intelligent Cloud Services.

Support for Microsoft Fabric OneLake Connector as a source in file listeners

File listener now supports Microsoft Fabric OneLake connector as a source.

For more information about configuring a file listener with a Microsoft Fabric OneLake connector as a source, see *Components*.

Taskflows

Taskflows include the following new features and enhancements:

- You can use the unpublish resource to unpublish multiple taskflows simultaneously and save time. To unpublish multiple taskflows simultaneously, use a POST request with the following URL:

```
<Informatica Intelligent Cloud Services URL>/active-bpel/asset/v1/unpublish
```

For more information about using the unpublish resource, see *Taskflows*.
- When you use an authenticated proxy server to run the ParamSetCli commands, you can use an encrypted password for the PROXYPWD parameter value in the `restenv.properties` file. For more information about using an encrypted proxy password for proxy server authentication in the ParamSetCli utility, see *REST API Reference*.

Changed behavior

Informatica Intelligent Cloud Services Data Integration includes the following changed behaviors in this release.

Importing assets with persisted values

When you import a mapping task that contains an in-out parameter or a non-shared sequence generator and the task does not already exist in the target location, by default, Informatica Intelligent Cloud Services resets the persisted value to the default parameter value or initial sequence value.

You can choose to retain persisted values when you import mapping tasks from the Informatica Intelligent Cloud Services UI. You cannot alter the default behavior when you import mapping tasks with the REST API.

Previously, Informatica Intelligent Cloud Services retained the persisted value in the export file by default.

For more information about importing assets, see *Asset Management*.

Connectors

The February 2024 release includes the following enhanced connectors.

Enhanced connectors

This release includes enhancements to the following connectors.

Databricks Delta Connector

This release includes the following enhancements for Databricks Delta Connector:

- You can use the Databricks cluster version 13.3 LTS at runtime to process Databricks Delta mappings and tasks.
- You can use key range partitioning to distribute the rows of source data based on the partition key and the range of values for each partition.

Google BigQuery V2 Connector

This release includes the following enhancements for Google BigQuery V2 Connector:

- The Google BigQuery V2 Connector uses Spark BigQuery connector library to write the data in a mapping in advanced mode. This library utilizes the Google BigQuery Storage API to directly write the data into Google BigQuery and does not use Google Cloud Storage for intermediate staging.
- You can configure standard and materialized views in the following scenarios within a mapping:
 - Configure the connection to use Google BigQuery Storage API for staging.
 - Configure the retry strategy in the connection.
 - Configure a source object with Record data type and repeated columns.
 - Configure the EXPORT DATA statement for staging.
 - Read the data in an advanced mode mapping.

Snowflake Data Cloud Connector

You can push a few additional functions such as bitwise expression, context, and aggregate functions to Snowflake when you create a mapping in SQL ELT mode.

For the list of supported functions, see the Snowflake Data Cloud Connector documentation.

Changed behavior

This release includes changes in behavior for the following connectors.

Google BigQuery V2 Connector

Effective in this release, Google BigQuery V2 Connector includes the following change in behavior:

- When you configure a mapping, you cannot use view or materialized view as a target object. Previously, you could import view or materialized view as a target object.
- If you have an existing mapping that writes a view as the target object, override the target object with the **Target Table Name** advanced property, and set the optional property `DisableMappingDeployment:true` at the Google BigQuery V2 connection, the mapping fails with the following error:

```
Operation failed: Internal Error Occurred. Contact Support : [Cannot create write operation. The node supports read operation only].
```

Previously, the mapping ran successfully.

To revert to the previous behavior, remove the optional property `DisableMappingDeployment:true` at the Google BigQuery V2 connection and run the mapping again.

- When you configure a mapping in advanced mode that writes the data of more than 10 MB to a row, the mapping fails.

Previously, the mapping ran successfully.

To revert to the previous behavior, contact Informatica Global Customer Support.

- When you configure a mapping in advanced mode and write Date, Datetime, and Timestamp values, `0001-01-01 00:00:00` or `0001-01-02` into a Timestamp column in the target, the mapping fails. Also, data corruption might occur when you write Date, Datetime, and Timestamp values less than `1910-01-01 00:00:00`.

Previously, the mapping ran successfully.

To revert to the previous behavior, contact Informatica Global Customer Support.

- When you configure a mapping in advanced mode with Write truncate selected in the **Write Disposition** target advanced property, and the target table contains Date or Datetime columns, the mapping fails. Previously, the mapping ran successfully.

To revert to the previous behavior, select Write append in the **Write Disposition** target advanced property and enable the **Truncate Target Table** target advanced property.

Microsoft Azure Synapse SQL Connector

Effective in this release, when you use the copy command to load data from the staging location to Microsoft Azure Synapse SQL and the data contains Windows newline characters, the Secure Agent writes the newline characters with both the carriage return (`/r`) and line feed (`/n`) to the target.

Previously, the Secure Agent wrote only the line feed (`/n`) to the target causing a change in the length of characters in the target data.

Snowflake Data Cloud Connector

Effective in this release, Snowflake Data Cloud Connector includes the following changes in behavior:

- When you run a mapping task to insert data from a CDC source to a Snowflake target where a primary key is not defined, the mapping task inserts data directly to the target table instead of creating a temporary stage table.

Previously, if the Snowflake target did not have a defined primary key, the mapping task inserted data from a CDC source to a temporary stage table before inserting to the Snowflake target.

- When you run a mapping task to read from a CDC source that does not contain any changed data and write to a Snowflake target for which the recovery mechanism is disabled, the mapping task does not commit data to the target.

Previously, even if a CDC source does not contain any changed data and the recovery mechanism is disabled for a Snowflake target, the mapping task executed the commit request and showed incorrect result.

CHAPTER 3

December 2023

The following topics provide information about new features, enhancements, and behavior changes in the December 2023 release of Informatica Intelligent Cloud ServicesSM Data Integration.

Connectors

The December 2023 release includes the following enhanced connectors.

Enhanced connectors

This release includes enhancements to the following connectors.

[Databricks Delta Connector](#)

You can push the IIF() function to process in Databricks Delta.

For the list of functions that you can push to Databricks Delta, see the [Databricks Delta Connector documentation](#).

CHAPTER 4

November 2023

The following topics provide information about new features, enhancements, and behavior changes in the November 2023 release of Informatica Intelligent Cloud ServicesSM Data Integration.

Important notices

The November 2023 release of Informatica Intelligent Cloud Services Data Integration includes the following important notices:

[Preview initiated](#)

Effective in the November 2023 release, the following functionality is available for preview:

Mappings in SQL ELT mode

Create a mapping in SQL ELT mode when you want the mapping logic to be processed by your cloud ecosystem.

When you run a mapping in SQL ELT mode, Data Integration translates the transformation logic into ecosystem-specific SQL statements and commands that run in the underlying cloud data warehouse. This increases the data processing speed because the data isn't moved out of the cloud infrastructure for processing. It also makes the data integration pipelines more efficient.

When you create a mapping in SQL ELT mode, you can read from and write to the following sources and targets:

- You can read data from your Snowflake data warehouse and write it to the same cloud data warehouse.
- You can read data from a Microsoft Azure Data Lake Storage Gen2 or Amazon S3 data lake and write it to your Snowflake data warehouse.

For more information about creating mappings in SQL ELT mode, see *Mappings*.

Note: This functionality is not enabled by default. If you'd like to try out this feature, contact Informatica Global Customer Support.

Preview functionality is supported for evaluation purposes but is unwarranted and is not supported in production environments or any environment that you plan to push to production. Informatica intends to include the preview functionality in an upcoming release for production use, but might choose not to in accordance with changing market or technical circumstances. For more information, contact Informatica Global Customer Support.

New features and enhancements

Informatica Intelligent Cloud Services Data Integration includes the following new features and enhancements in this release.

CLAIRE-powered runtime strategies

CLAIRE-powered runtime strategies include the following new features and enhancements:

- Mapping tasks use a new layout to display runtime strategies. The option to enable or disable runtime strategies was replaced by an option to run the task using the default runtime strategy.
- CLAIRE can recommend runtime strategies for mapping tasks that run in Databricks Delta ecosystems. The mapping task must be based on a mapping in advanced mode.

For more information, see *Tasks*.

Expression transformation field types

When you have a long list of types and want to see the type for each expression field, add the **Field Type** column in the Expression properties table.

Generate a single output in the Hierarchy Builder transformation

You can configure a Hierarchy Builder transformation to group all of the output elements in a single output group.

Start REST API session using JWT access token

You can start a new REST API session by logging in using a JWT access token for authentication. To log in, use the loginOAuth v2 REST API resource and include a token from the identity provider in the request.

For more information, see REST API Reference.

Update dynamic cache condition

When a Lookup transformation uses a dynamic lookup cache, you can configure a conditional expression to tell Data Integration when to update the cache. Data Integration updates the cache when the expression is true.

For more information about configuring Lookup transformations, see *Transformations*.

User-defined functions supported in advanced mode

You can include user-defined functions in the following transformations in advanced mode:

- Aggregator
- Expression

Availability of human task assets in Data Integration Bundles

When you create a bundle in Data Integration, you can include human task assets. The users can use the bundles for various Informatica Intelligent Cloud Services projects.

Taskflows

Taskflows include the following new features and enhancements:

- You can use an API to terminate up to 200 taskflows. Use the following URI to terminate multiple taskflows:

```
PUT <Informatica Intelligent Cloud Services URL>/active-bpel/services/tf/terminate
```

Enter the taskflow run IDs as the input in the body of the PUT request. The API supports basic authentication and session ID authorization.

For more information about terminating multiple taskflows, see *Taskflows*.

- You can use a JLS API to get a list of up to 200 taskflows that are suspended. Use the following URI to get the list of suspended taskflows:

```
{{<Informatica Intelligent Cloud Services URL>}}/jls-di/api/v1/Orgs(<Organization ID>)/JobLogEntries?$top=<Number of suspended taskflows>&$filter=(status eq 'SUSPENDED' and assetType eq 'TASKFLOW')
```

Use the following URI to get the list of the latest suspended taskflows in an organization filtered by user:

```
{{<Informatica Intelligent Cloud Services URL>}}/jls-di/api/v1/Orgs(<Organization ID>)/JobLogEntries?$top=<Number of suspended taskflows>&$filter=(status eq 'SUSPENDED' and startedBy eq '<username>' and assetType eq 'TASKFLOW')
```

By default, the number of suspended taskflows that the API returns is 200.

For more information about getting the list of suspended taskflows, see *Taskflows*.

- You can add a custom name to a taskflow by using an API and the RunAJob utility. You can use the custom name as an identifier for the taskflow.
For more information about adding a custom name to a taskflow, see *Taskflows*.
- You can view the metering information for the taskflows that are run within your organization and sub-organizations on the **Metering** page of Administrator. You can also download metering usage data from the taskflow metering page. The metering page includes details such as the date when the taskflows within the organization were run, organization ID, organization type, environment name, environment type, tier, IPU per unit, total number of taskflows that were run, and the number of IPU's consumed.
For more information about taskflow metering, see *Organization Administrator*.

Changed behavior

Informatica Intelligent Cloud Services Data Integration includes the following changed behaviors in this release.

Partitions in data flows

When you configure a mapping with multiple data flows, sources in different flows can have a different number of partitions.

Previously, sources in different flows had to have the same number of partitions.

Platform REST API

If you send a pull request and it fails for any reason after initialization, you can find the pull status and error in the source control logs or by sending a sourceControlAction request.

Previously, if you sent a pull request and the object couldn't be found in the repository, the pull operation returned the pull status and error in the response body and excluded it from the source control logs. All other failures were visible in the source control logs.

Pushdown optimization is now SQL ELT optimization

Pushdown optimization has been renamed to SQL ELT optimization in the user interface and documentation.

The user interface includes this change in many places, including the following instances:

- In the Mapping Designer, the Pushdown Optimization panel is now called the SQL ELT Optimization panel. Use the SQL ELT Optimization panel to perform SQL ELT optimization preview and SQL ELT optimization data preview.
- In the mapping task wizard and the mapping task editor, the **Pushdown Optimization** property is now called **SQL ELT Optimization**.
- When you use CLAIRE-powered runtime strategies, CLAIRE might recommend using SQL ELT optimization to run the mapping.
- In the job details for a mapping task, the Pushdown Optimization area is now called SQL ELT Optimization.

For more information about SQL ELT optimization, see *Tasks*. For more information about SQL ELT optimization preview, see *Mappings*. For more information about monitoring mapping tasks, see the Monitor help.

System-generated keys in the Hierarchy Parser transformation

The Hierarchy Parser transformation always generates primary and foreign keys sequentially starting from 1.

Previously, the transformation didn't generate keys in sequential order in some cases.

Taskflows

On the **Advanced** tab, the **Disabled concurrent execution** option to disable the concurrent taskflow execution is renamed to **Skip if Running**.

Connectors

The November 2023 release includes the following enhanced connectors.

Enhanced connectors

This release includes enhancements to the following connectors.

Amazon Redshift V2 Connector

This release includes the following enhancements for Amazon Redshift V2 Connector:

- You can read and write data of the SUPER data type.
- You can configure the IS NULL and IS NOT NULL operators in simple filters and uncached lookups.
- You can push REG_MATCH and ASCII functions to an Amazon Redshift database when you configure full SQL ELT optimization.
- You can limit the number of lookups using the lookup data filter condition in a lookup source from a mapping that reads data from Amazon Redshift. This feature does not apply to mappings in advanced mode.

Amazon S3 V2 Connector

You can configure folder-level access to read data from and write data to particular folders in an Amazon S3 bucket.

Business 360 FEP Connector

This release includes the following enhancements for Business 360 FEP Connector:

- You can use Business 360 FEP Connector to securely read data from the MDM SaaS data store.
- You can now update field values and delete field group entries in record sections. To update field values or delete field group entries in record sections, you can use the **Update** or **Delete** operation in a target transformation object that uses a Business 360 FEP connection.

Databricks Delta Connector

This release includes the following enhancements for Databricks Delta Connector:

- **SQL ELT enhancements**
You can push additional functions to process in Databricks Delta.
For the list of functions that you can push to Databricks Delta, see the Databricks Delta Connector documentation.
- You can use the JDBC URL versions 2.6.25 or later to connect to SQL warehouse.
- You can use the personal staging location on Windows to stage data locally for mappings and tasks.
- When you write data to multiple targets in a single pipeline, you can configure the Secure Agent and task properties to process the target operations in a specific order.
- You can limit the number of lookups using the lookup data filter condition in a lookup source from a mapping that reads data from Databricks Delta. This feature does not apply to mappings in advanced mode.

Flat file Connector

You can specify a date format in the advanced source and target properties in a mapping.

Google BigQuery V2 Connector

This release includes the following enhancements for Google BigQuery V2 Connector:

- You can push DATE_DIFF() and MD5() functions to the Google BigQuery database when you configure SQL ELT optimization.
- You can enable connection retry at Google BigQuery connection when you read data in direct mode.
- You can use standard views for the source objects.
- You can limit the number of lookups using the lookup data filter condition in a lookup source from a mapping that reads data from Google BigQuery. This feature does not apply to mappings in advanced mode.

Kafka Connector

The Kafka connector supports header and offset fields at the source.

Microsoft Azure Synapse SQL Connector

This release includes the following enhancements for Microsoft Azure Synapse SQL Connector:

- You can configure Service Principal authentication to connect to Microsoft Azure Synapse SQL using the application ID or client ID for your application registered in Azure Active Directory.

- You can limit the number of lookups using the lookup data filter condition in a lookup source from a mapping that reads data from Microsoft Azure Synapse SQL. This feature does not apply to mappings in advanced mode.

SAP Table Connector

SAP Table Connector uses the latest transport requests ER6K900059 for the TABLE_READER transport and S4NK900085 for the TABLE_READER_Addon transport.

Snowflake Data Cloud Connector

This release includes the following enhancement for Snowflake Data Cloud Connector:

- You can create a mapping in SQL ELT mode when you extract data from the following sources and load data to Snowflake:
 - Amazon S3
 - Microsoft Azure Data Lake Storage Gen2
 - Snowflake
- You can limit the number of lookups using the lookup data filter condition in a lookup source from a mapping that reads data from Snowflake. This feature does not apply to mappings in advanced mode.

Changed behavior

This release includes changes in behavior for the following connectors.

Pushdown optimization is now SQL ELT optimization

Pushdown optimization has been renamed to SQL ELT optimization in the user interface and documentation.

The user interface includes this change in many places, including the following instances:

- In the Mapping Designer, the Pushdown Optimization panel is now called the SQL ELT Optimization panel. Use the SQL ELT Optimization panel to perform SQL ELT optimization preview and SQL ELT optimization data preview.
- In the mapping task wizard and the mapping task editor, the **Pushdown Optimization** property is now called **SQL ELT Optimization**.
- When you use CLAIRE-powered runtime strategies, CLAIRE might recommend using SQL ELT optimization to run the mapping.
- In the job details for a mapping task, the Pushdown Optimization area is now called SQL ELT Optimization.

For more information about SQL ELT optimization, see the help for the appropriate connector.

Amazon Redshift V2 Connector

Effective in this release, Amazon Redshift V2 Connector includes the following change in behavior:

- If you configure ASCII() and REG_MATCH() operators in a mapping enabled with SQL ELT optimization to read data that contains Null and Empty values, the mapping treats Null as NULL and Empty as Empty. Previously, the mapping ran without SQL ELT optimization. Data Integration treated these values as either Null or Empty depending on whether you selected the Treat NULL as NULL property in the Amazon Redshift Source transformation properties.
- When you run a mapping enabled for SQL ELT optimization and the mapping contains multiple pipelines and one of the targets in the intermediate pipeline fails, the mapping runs without SQL ELT optimization. Previously, if one of the targets in the intermediate pipeline fails, the mapping ran with partial SQL ELT optimization.

Business 360 FEP Connector

Effective in this release, when you try to configure operation type for a target transformation, you can no longer view the **Data Driven** operation.

Previously, a Business 360 FEP connection listed the **Data Driven** operation but used the upsert operation when you select it.

Google BigQuery V2 Connector

Effective in this release, Google BigQuery V2 Connector includes the following change in behavior:

- After the upgrade, a mapping enabled with full SQL ELT optimization and configured with the DATE_DIFF() function will run with full SQL ELT optimization.
Previously, a mapping enabled with SQL ELT optimization and configured with the DATE_DIFF() function ran with source SQL ELT optimization or without SQL ELT optimization.
- When you configure the DATE_DIFF function in an Expression transformation and the transformation output is set to double or decimal data type in a mapping enabled with SQL ELT optimization, the function returns an integer value. However, when you run the mapping without SQL ELT optimization, the function returns a double value. Also, you might encounter a data mismatch when the transformation output is set to integer data type.
To revert to the previous behavior, set the optional property `SkipFunctionsInPDO:DATE_DIFF` at the Google BigQuery connection.

Microsoft Azure Data Lake Storage Gen2 Connector

Effective in this release, a mapping configured to create a new Microsoft Azure Data Lake Storage Gen2 target at runtime, along with the append strategy, creates a target file and writes the data to it. The mapping appends data in subsequent runs.

Previously, the mapping failed to create a new target when you configured the append strategy and resulted in an error.

Microsoft Azure Synapse SQL Connector

Effective in this release, when you push a CONCAT() function for data types other than string to process in Microsoft Azure Synapse SQL from a mapping task enabled with SQL ELT optimization, the task runs successfully with SQL ELT optimization.

Previously, the mapping ran without SQL ELT optimization.

CHAPTER 5

Upgrade

The following topics provide information about tasks that you might need to perform before or after an upgrade of Informatica Intelligent Cloud Services Data Integration. Post-upgrade tasks for previous monthly releases are also included in case you haven't performed these tasks after the previous upgrade.

Preparing for the upgrade

Perform the following tasks to prepare for upgrade to the April 2024 release.

Disk space requirements and file preparation

The Secure Agent upgrades the first time that you access Informatica Intelligent Cloud Services after the upgrade.

Files that you added to the following directory are preserved after the upgrade:

```
<Secure Agent installation directory>/apps/Data_Integration_Server/ext/deploy_to_main/bin/  
rdtm-extra
```

Note: After a major release upgrade, the old and new versions of the services might continue to run concurrently for up to 24 hours. Ensure that the Secure Agent machine has sufficient free memory to avoid memory depletion during this period.

Perform the following steps to ensure that the Secure Agent is ready for the upgrade:

1. Ensure that each Secure Agent machine has sufficient disk space available for upgrade.

The machine must have at least 5 GB free space or the amount of disk space calculated using the following formula, whichever is greater:

```
Minimum required free space = 3 * (size of current Secure Agent installation directory -  
space used for logs directory)
```

2. Close all applications and open files to avoid file lock issues, for example:

- Windows Explorer
- Notepad
- Windows Command Processor (`cmd.exe`)

Dynamic files in preprocessing commands

Data Integration refreshes mapping and task metadata during an upgrade. If you have a task that uses preprocessing commands to create a dynamic file, ensure that a file with the referenced file name is present at the specified file path before the first task run after the upgrade.

If a file isn't present, Data Integration can't map the metadata between the preprocessing command and the referenced file, and the task fails. After the first task run, Data Integration retains the file path.

Post-upgrade tasks for the February 2024 release

Perform the following tasks after your organization is upgraded to the February 2024 release.

Hive Connector

The Informatica Hive third-party script and the Informatica Hive third-party properties files for Hive Connector include updates in the third-party JAR names for the Azure HDInsight 4.0 distribution.

After the upgrade, if you are using the Azure HDInsight 4.0 distribution for Hive Connector, your existing Hive mappings fail.

Before you run existing mappings, run the Informatica Hive third-party script and download the required Azure HDInsight 4.0 distribution files that are available as part of the Hive Connector package in the Secure Agent installation directory. Then, restart the Secure Agent.

For more information about the tasks that you need to perform, see the *Hive Connector* documentation.

INDEX

C

Cloud Application Integration community
URL [5](#)
Cloud Developer community
URL [5](#)

D

Data Integration community
URL [5](#)

I

Informatica Global Customer Support
contact information [6](#)
Informatica Intelligent Cloud Services
web site [5](#)

M

maintenance outages [6](#)

S

Secure Agents
upgrade preparation [32](#)
status
Informatica Intelligent Cloud Services [6](#)
system status [6](#)

T

trust site
description [6](#)

U

upgrade notifications [6](#)
upgrade preparation
Secure Agent preparation [32](#)

W

web site [5](#)