



Informatica® Data Ingestion and Replication  
October 2024

# Data Ingestion and Replication What's New

Informatica Data Ingestion and Replication Data Ingestion and Replication What's New  
October 2024

© Copyright Informatica LLC 2019, 2025

Publication Date: 2025-07-03

# Table of Contents

- Chapter 1: October 2024..... 4**
  - Important notices. . . . . 4
  - New features and enhancements. . . . . 4
    - Common. . . . . 5
    - Database Ingestion and Replication. . . . . 5
  - Changed behavior. . . . . 6
- Chapter 2: August 2024..... 7**
  - Important notices. . . . . 7
  - New features and enhancements. . . . . 7
    - Common. . . . . 7
    - Application Ingestion and Replication. . . . . 8
    - Database Ingestion and Replication. . . . . 8
  - Changed behavior. . . . . 9
  - Post-upgrade task for the August 2024 release. . . . . 9
- Chapter 3: July 2024..... 11**
  - New features and enhancements. . . . . 11
    - Common. . . . . 11
    - Application Ingestion and Replication . . . . . 11
    - Database Ingestion and Replication. . . . . 12
    - Streaming Ingestion and Replication. . . . . 12
  - Changed behavior. . . . . 12
  - Post-upgrade tasks for the July 2024 release. . . . . 13
- Chapter 4: May 2024..... 14**
  - New features and enhancements. . . . . 14
    - Common. . . . . 14
    - Mass Ingestion Databases. . . . . 16
  - Changed behavior. . . . . 17
- Index..... 18**

# CHAPTER 1

## October 2024

This section provides information about new features, enhancements, and behavior changes in the October 2024 release of Data Ingestion and Replication.

### Important notices

The October 2024 release includes the following important notice.

#### [New wizard for configuring application and database ingestion and replication tasks](#)

A new wizard for configuring application ingestion and replication tasks and database ingestion and replication tasks is available to new users and to selected organizations on a limited basis. The wizard is designed to provide a simpler, cleaner user interface.

The new interface includes the following benefits:

- You can define a primary cloud data warehouse destination once and have it appear by default when you define a task from the new wizard. At initial release, the new wizard supports only Snowflake as the primary cloud data warehouse destination.
- You have the option of defining new source and target connections within the wizard flow. The wizard provides embedded help to help you set connection properties.
- Optional advanced source and target properties are initially hidden so that you can focus on the required fields first. You can easily expose the advanced properties with a single click.

To get access to the new wizard, contact Informatica Global Customer Support or your Customer Success Manager. An Informatica representative will then contact you with instructions for accessing the new interface.

For a brief tour of the new interface, watch [this video](#).

For access to the documentation for new task configuration wizard, go to [this KB](#).

### New features and enhancements

The October 2024 release of Data Ingestion and Replication includes the following new features and enhancements.

## Common

The October 2024 release of Data Ingestion and Replication includes the following new features that are common to multiple types of ingestion and replication tasks.

### New job metrics command available in the Data Ingestion and Replication Command-Line Interface

The Data Ingestion and Replication Command-Line Interface (CLI) provides a new `job metrics` command to display statistics for a specified job or a subtask of a job. This enhancement is available for application ingestion and replication jobs and database ingestion and replication jobs.

On the job level, the statistics include records read, records written, and capture progress.

On the subtask level, you can display the following statistics:

- Records read and records written for initial load jobs
- Counts of processed inserts, updates, deletes, and LOB changes for incremental and combined load jobs
- The number of records processed and replicated during the initial load phase of combined load jobs

You can use the `-f` field parameter to specify which of these values should be included in the output.

For more information, see the *Data Ingestion and Replication Command-Line Interface*.

### Support for Kerberos authentication for Oracle sources and targets

Kerberos authentication is supported for connecting to Oracle sources and targets in database ingestion and replication tasks that use any load type. For CDC sources, it's supported for either the Query-based or Log-based CDC method. It's also supported for connecting to Oracle targets in application ingestion and replication tasks.

When you define an Oracle Database Ingestion connection, select **Kerberos** in the **Authentication Mode** field. You must also perform some prerequisite Kerberos configuration tasks.

For more information, see the Oracle Database Ingestion connection in *Connectors and Connections*.

## Database Ingestion and Replication

The October 2024 release of Database Ingestion and Replication includes the following new features and enhancements:

### Soft deletes for Google BigQuery targets

Database ingestion and replication incremental load jobs and combined initial and incremental load jobs can now process DML delete operations on the source as soft deletes on Google BigQuery targets. Database Ingestion and Replication marks the soft-deleted records with a "D" in the `INFA_OPERATION_TYPE` column on the target without actually deleting the records.

For more information, see *Database Ingestion and Replication* > Configuring a database ingestion and replication task > Configuring the target > Google BigQuery target properties.

### Support for multiple Oracle log locations for jobs that use BFILE access with directory objects

If you use BFILE access to Oracle data and the data is stored in different log locations, Database Ingestion and Replication can now automatically find the logs by querying for all directory objects that begin with `ARCHIVELOG_DIR` and `ONLINELOG_DIR` and end with any *suffix*, for example, `ARCHIVELOG_DIR_01`, `ARCHIVE_LOG_DIR_02`, and so on. This feature enables Database Ingestion and Replication to support Oracle standby databases that use log locations different from those on the primary database, replica logs on Amazon RDS database instances, and archive logs in the Fast Recovery Area to which the

USE\_DB\_RECOVERY\_FILE\_DEST parameter points. You can override the ARCHIVELOG\_DIR and ONLINELOG\_DIR names by using the custom properties pwx.cdcreader.oracle.database.additionalBFILEARCHIVEDIR=<directory\_object> and pwx.cdcreader.oracle.database.additionalBFILEONLINEDIR=<directory\_object>.

If you grant the CREATE ANY DIRECTORY and DROP ANY DIRECTORY privileges to database users who run database ingestion and replication jobs, the directory objects can be created at runtime, as needed. Database Ingestion and Replication performs no cleanup processing on these generated directory objects.

If you use the USE\_DB\_RECOVERY\_FILE\_DEST parameter and the CREATE ANY DIRECTORY and DROP ANY DIRECTORY privileges have *not* been granted, your DBA must create the directory objects daily or weekly, before your database ingestion and replication jobs run, by using a script such as:

```
create or replace directory ARCHIVELOG_DIR_2024_08_19 as '<DB_RECOVERY_FILE_DEST>/
2024_08_19'
```

If you use the USE\_DB\_RECOVERY\_FILE\_DEST parameter and the database user has the CREATE ANY DIRECTORY and DROP ANY DIRECTORY privileges, the directory objects are created as needed at runtime and dropped after 14 days. These directory objects have the naming convention <ARCHIVEDIRNAME>\_YYYY\_MM\_DD.

For more information, see *Database Ingestion and Replication* > Database Ingestion and Replication sources > Oracle sources.

## Changed behavior

The October 2024 release of Data Ingestion and Replication includes the following changed behaviors.

### Optimizing use of source unique indexes when no primary key exists

After the October 2024 upgrade, new application ingestion and replication tasks and database ingestion and replications tasks preferentially use the source primary key when generating Oracle target objects. If a source does not have a primary key but does have unique indexes, the job chooses the best unique index for replication in the following order of priority: 1) the unique index that has only NOT NULL columns, 2) if multiple unique indexes have only NOT NULL columns, the unique index with the least number of columns, or 3) if no source primary key or unique NOT NULL index exists, the unique index with the greatest number of columns.

## CHAPTER 2

# August 2024

This section provides information about new features, enhancements, and behavior changes in the August 2024 release of Data Ingestion and Replication.

## Important notices

The August 2024 release includes the following important notices.

### [Preview lifted](#)

Effective in the August 2024 release, the ability to share a Secure Agent group that includes multiple agents with sub-organizations is no longer in preview. It's now production ready for Application Ingestion and Replication users and Database Ingestion and Replication users. The administrator of the parent organization can share the parent organization's Secure Agent group with one or more sub-organizations. When sub-organization members create an ingestion and replication task of any load type, they can select the shared Secure Agent group as the runtime environment. Then when they run their ingestion and replication jobs, the jobs run on a Secure Agent in the group. If the Secure Agent becomes unavailable, the job can switch to another Secure Agent in the group. All of the agents in a Secure Agent group must use the Database Ingestion agent service.

## New features and enhancements

Read about new features and enhancements in the August 2024 Data Ingestion and Replication release.

### Common

The August 2024 release of Data Ingestion and Replication includes the following new features that are common to multiple types of ingestion and replication tasks.

### [SSL encryption support for SQL Server target connections](#)

Application ingestion and replication jobs and database ingestion and replication jobs that have a SQL Server target type (including Azure SQL Database, RDS for SQL Server, and Azure SQL Managed Instance) can use SSL encryption for the target connection. This feature is supported for jobs with any supported source type and load type. The connection's authentication mode must be either SQL Server Authentication or Windows Authentication v2.

When you configure SQL Server connection properties, set the **Encryption Method** property to **SSL**, **Request SSL**, or **Login SSL**. Then set the related properties, including Crypto Protocol Version, Validate Server Certificate=True, Trust Store, Trust Store Password, and Host Name in Certificate.

For more information, see Connectors and Connections > *Microsoft SQL Server connection properties*.

### [New alert condition for incremental load and combined load jobs not processing DML changes](#)

For application ingestion and replication jobs and database ingestion and replication jobs, you can now configure alert email notifications if an incremental load or a combined initial and incremental load job has not been processing any DML data for a significant period of time.

You can create the alert notification from the **Alerts > Data Ingestion and Replication Alerts** page in Operational Insights. In the **Create Alert** wizard, select the **Source CDC Delay** option to receive notifications if a delay in processing records from the source occurs.

## Application Ingestion and Replication

The August 2024 release of Application Ingestion and Replication includes the following new feature and enhancement:

### [Support for Microsoft Dynamics 365 sources with Microsoft Azure SQL Database targets for all load types](#)

Application Ingestion and Replication supports Microsoft Dynamics 365 sources with Microsoft Azure SQL Database targets for initial, incremental, and combined load types.

For more information, see *Guidelines for Microsoft Azure SQL Database targets*.

## Database Ingestion and Replication

The August 2024 release of Database Ingestion and Replication includes the following new features and enhancements:

### [Active Directory Password authentication for SQL Server sources](#)

You can now use Active Directory Password authentication for SQL Server source connections that are used by database ingestion and replication initial load jobs and by incremental load and combined load jobs that use the **CDC Tables** or **Log-based** CDC method. This authentication method is not currently supported for jobs that use the **Query-based** CDC method.

For more information, see Connectors and Connections > *Microsoft SQL Server connection properties*.

### [Replacement of special characters in column names with a provided choice of string in the Snowflake target](#)

If the source tables contain columns that have special characters in the name, you can now replace the special characters with a provided choice of string in the Snowflake target using the `targetReplacementValue=<string>` custom property.

To replace the special characters with an underscore, set the `targetReplacementValue=<string>` custom property on the **Target** page of the task wizard and in a properties file which you can define in the **Custom Configuration Details** area of the Secure Agent.

For more information, see Database Ingestion and Replication > *Snowflake targets*.



# Changed behavior

The August 2024 release of Data Ingestion and Replication includes the following changed behaviors.

## Enhanced CDC processing of persisted computed columns in SQL Server sources

If you use the **CDC Tables** method to capture change data from SQL Server sources that include persisted computed columns, database ingestion and replication incremental load and combined load jobs now replicate the computed column expression values instead of nulls. If you want to revert to the prior behavior of replicating nulls for persisted computed columns, set the `pxw.custom.sslr_cdcread_replicate_persisted_cols` custom property to false on the **Source** page of the task wizard.

## Active management of SQL Server CDC tables by default

For database ingestion and replication incremental load jobs that have SQL Server sources and capture data from CDC tables, which primarily occurs when using the **CDC Tables** capture method, Database Ingestion and Replication now uses active management of the CDC tables by default. Activate management enables enhanced processing of source DDL and DML changes. Previously, to use active management of CDC tables, you had to set the `pxw.custom.sslr_cdc_manage_instances` source custom property to true. Now the default value for this property is true.

## Newly added SAP HANA source columns are ignored in incremental load jobs

For database ingestion and replication incremental load jobs that have a SAP HANA source, when a new column is added to the source table and it is not present in the shadow \_CDC tables, the new column is ignored.

Previously, when a new column of a supported SAP HANA data type was added to the source table, it was included in the source definition and the job failed because a corresponding column was not found in the shadow \_CDC tables. If you want to revert to the prior behavior, set the `pxw.cdcreader.HANA.addCollIgnore` custom property to false on the **Source** page of the task wizard.

## Support for Avro logical types

Newly deployed application ingestion and replication jobs and database ingestion and replication jobs that use the Avro output format now map source types to the Avro logical types of date, decimal, and timestamp (with microsecond precision) by default. For jobs that have already been deployed or that are redeployed, the prior behavior of mapping all source types to the Avro primitive type of string by default continues. The change to using Avro logical types helps downstream applications and processes process the Avro data correctly. If you want to revert to the prior behavior, you can set the custom property `avroEnableLogicalTypes` to false on the **Target** page of the task wizard.

# Post-upgrade task for the August 2024 release

Perform the following task after Data Ingestion and Replication is upgraded to the August 2024 release.

## Upgrade consideration for Databricks using third-party JAR files

After the August 2024 upgrade, Application Ingestion and Replication and Database Ingestion and Replication include the Databricks JDBC and Simba Spark JDBC drivers in the Secure Agent package.

The package contains the following JDBC driver versions:

- Databricks JDBC driver: 2.6.36

- Simba Spark JDBC driver: 2.6.21.1039

To ensure consistency Informatica recommends that you delete the existing version of these JAR files from the following directory:

```
\Database_Ingestion\ext
```

The existing application ingestion and replication tasks and database ingestion and replication tasks deployed before the upgrade will continue to run without issues when you resume these tasks after the upgrade.

# CHAPTER 3

## July 2024

This section provides information about new features, enhancements, and behavior changes in the July 2024 release of Data Ingestion and Replication.

### New features and enhancements

Read about new features and enhancements in the October 2024 Data Ingestion and Replication release.

#### Common

The October 2024 release of Data Ingestion and Replication includes the following new features that are common to multiple types of ingestion and replication tasks.

##### [New job metrics command available in the Data Ingestion and Replication Command-Line Interface](#)

The Data Ingestion and Replication Command-Line Interface (CLI) provides a new `job metrics` command to display statistics for a specified job or a subtask of a job. This enhancement is available for application ingestion and replication jobs and database ingestion and replication jobs.

On the job level, the statistics include records read, records written, and capture progress.

On the subtask level, you can display the following statistics:

- Records read and records written for initial load jobs
- Counts of processed inserts, updates, deletes, and LOB changes for incremental and combined load jobs
- The number of records processed and replicated during the initial load phase of combined load jobs

You can use the `-f` field parameter to specify which of these values should be included in the output.

#### Application Ingestion and Replication

The July 2024 release of Application Ingestion and Replication includes the following new feature and enhancement:

##### [Schema drift support for PostgreSQL targets](#)

Application Ingestion and Replication adds support for automatic schema drift detection and handling for incremental load and combined initial and incremental load jobs that have a Salesforce source and a PostgreSQL target.

## Database Ingestion and Replication

The October 2024 release of Database Ingestion and Replication includes the following new features and enhancements:

### Soft Deletes for Google BigQuery targets

Database ingestion and replication incremental load jobs and combined initial and incremental load jobs can now process hard delete operations on the source as soft deletes on Google BigQuery targets. Database Ingestion and Replication marks the soft-deleted records with a "D" in the INFA\_OPERATION\_TYPE column on the target without actually deleting the records.

## Streaming Ingestion and Replication

The July 2024 release of Streaming Ingestion and Replication includes the following new feature and enhancement:

### Support for Business 360 Events Connector

Mass Ingestion Streaming now supports Business 360 Events connector as a source to transfer files.

## Changed behavior

The July 2024 release of Data Ingestion and Replication includes the following changed behaviors.

### Monitor now shows the Rows Processed count for ingestion and replication jobs

On the **Running Jobs** and **All Jobs** pages in Monitor and on the **My Jobs** page accessed from the Data Integration unified Home page, the **Rows Processed** column now displays values for running and successfully completed application ingestion and replication, database ingestion and replication, and streaming ingestion and replication jobs. For incremental load and combined initial and incremental load jobs, the count is displayed after CDC processing ends. The count is the number of rows written to the target. Previously, only "N/A" was displayed.

### Automatic switchover to another Secure Agent in a Secure Agent group

If the active Secure Agent on which application ingestion and replication jobs or database ingestion and replication jobs are running goes down unexpectedly, the jobs can automatically switch over to another available agent in the group after the 15-minute heartbeat interval elapses. You no longer have to Stop and Resume the job for the switchover to occur. Automatic switchovers are supported for:

- Application ingestion and replication and database ingestion and replication initial load jobs that have any source type and a target type other than Kafka
- Database ingestion and replication incremental load and combined load jobs that have an Oracle or SQL Server source, subject to the following limitations:
  - The job cannot have a Kafka target.
  - The job cannot have persistent storage enabled.
  - The job cannot use the Query-based CDC method to capture changes from the Oracle or SQL Server source.

# Post-upgrade tasks for the July 2024 release

Perform the following task after Data Ingestion and Replication is upgraded to the July 2024 release.

## Upgrade consideration for Google BigQuery using third-party jars

After the July 2024 upgrade, application ingestion and replication jobs and database ingestion and replication jobs might fail if you previously copied the Google BigQuery jar files, including jackson-databind-2.x.x.jar, jackson-core-2.x.x.jar, and jackson-annotations-2.x.x.jar, to the following location:

```
<INFA_HOME>/apps/Database_Ingestion/ext/
```

The jobs fail with the following error because the versions of the earlier jackson-databind-2.x.x.jar, jackson-core-2.x.x.jar, and jackson-annotations-2.x.x.jar files differ from the current 2.17.1 version of the .jar files that Application Ingestion and Replication and Database Ingestion and Replication use:

```
[DBMIP_20009] The process [CDC_COMBINED-SSLRCDReader-CDC_COMBINED PWX Reader(63)-main]
encountered an unexpected error. Error: [PWX session start failed with
StartSession({host=localhost,port=44991,sl=true,sslTrustStore=}) failed. Caused by:
Failed to open a socket. Error:Connection refused (Connection refused)].]
```

You must delete the older jackson-databind-2.x.x.jar, jackson-core-2.x.x.jar, and jackson-annotations-2.x.x.jar files and replace them with the current jackson-databind-2.17.1.jar, jackson-core-2.17.1.jar, and jackson-annotations-2.17.1.jar files from the following location:

```
<INFA_HOME>/downloads/package-GoogleBigQueryV2.{latestversion}/package/bigqueryv2/
thirdparty/infa.bigqueryv2
```

Then restart the Secure Agent and resume your jobs.

## Upgrade consideration for database ingestion and replication jobs that have Db2 for LUW sources

If a database ingestion and replication job that has a Db2 for LUW source tables with LOB columns was running during the upgrade to the July 2024 release and a Drop Column operation occurred on a source table, perform the following steps to ensure that the source and target definitions are still consistent in structure:

1. Stop the job.
2. Edit the associated task to select the **Include LOBs** option for the source.
3. Redeploy the job.

## CHAPTER 4

# May 2024

This section provides information about new features, enhancements, and behavior changes in the May 2024 release of Data Ingestion and Replication service.

## New features and enhancements

Read about new features and enhancements in the May 2024 Mass Ingestion release.

### Common

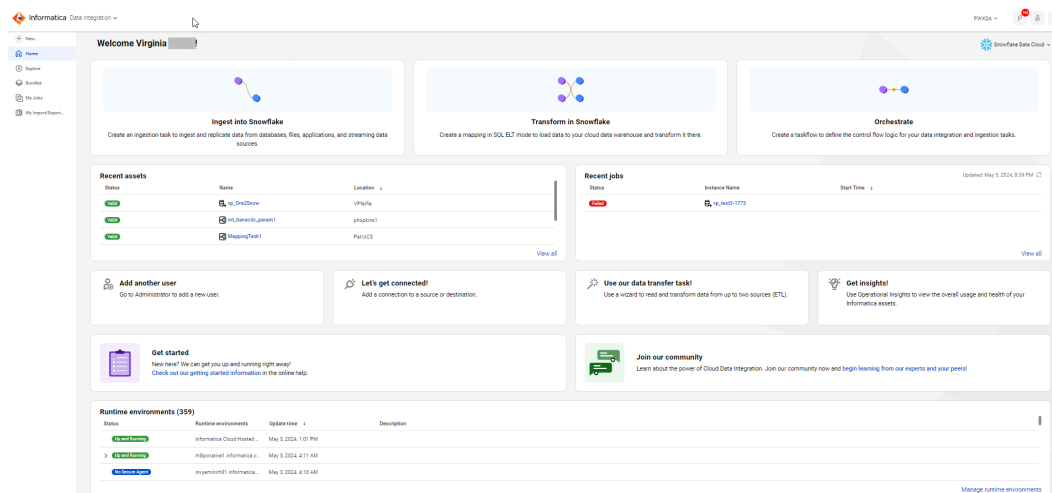
The May 2024 release of Informatica Intelligent Cloud Services Data Ingestion and Replication service includes the following new features that are common to multiple types of ingestion tasks.

#### Unified Home page

Most Mass Ingestion organizations will see a new Home page that provides a guided and unified experience for Data Integration and Data Ingestion and Replication users.

The Home page provides a central place from which you can create tasks, mappings, taskflows, and other assets, depending on your licensing. You can also access your recently created or updated assets and check the statuses of your recent jobs.

The panels that appear on the **Home** page are customized based on criteria such as the organization's licenses and the user's roles. For example, the following Home page is displayed for an Admin user in an organization that uses both Data Ingestion and Replication and Data Integration:



The navigation bar on the left side still provides the traditional New, Explore, and My Jobs options for creating a new task, exploring all assets, and viewing all of your jobs.

Limitations in the May 2024 release include:

- If you use the localized Japanese version of the user interface, or if you onboard to Data Ingestion and Replication after the May 2024 release is rolled out, the pre-existing Data Integration Home page appears instead of the new unified Home page. On this Home page, you can use the New, Explore, and My Jobs options in the navigation bar to create ingestion tasks, access your assets, and view details for your jobs. The new unified Home page is planned to be available to all Data Ingestion and Replication users in the July 2024 release.
- If you use the **Do you use a cloud data warehouse as your primary destination?** panel on the Home page to define a cloud data warehouse target like Snowflake, the information will not appear on the **Target** page of the task wizard when you create a new ingestion task. However, you can define the target from the task wizard.

## New UI look and feel

Along with the new unified home page, the Data Ingestion and Replication user interface has a fresh look and feel with a lighter palette, updated icons, and new fonts.

## Monitoring ingestion jobs in the Monitor service and from the unified Home page

You can now view and monitor ingestion jobs, along with other job types, in the Monitor service. You can monitor the progress and status of all jobs and all running jobs on the **All Jobs** and **Running Jobs** pages, respectively. On either page, you can click any job name to drill down to detailed information about the job. You can also perform some actions on the listed jobs, depending on the job status and task type, by using the **Actions** menu at the right end of each job row. To help you find your jobs more easily, the **All Jobs** page provides Filter, Find, and Sort options.

Alternatively, you can monitor all of your jobs, including ingestion and integration jobs, on the **My Jobs** page that is accessed by clicking **My Jobs** on the navigation bar of the unified Home page.

**Note:** You can continue to monitor only ingestion jobs on the **Mass Ingestion** page in Operational Insights.

## Microsoft Fabric OneLake targets in incremental load and combined load jobs

You can now use Microsoft Fabric OneLake as a target for application ingestion and database ingestion incremental load and combined initial and incremental load jobs.

Previously Microsoft Fabric OneLake was available as a target in initial load jobs only.

### Secure Agent identified in the application ingestion and database ingestion job details.

The name of the Secure Agent on which an application ingestion or database ingestion job runs is now displayed in the job-specific details. To access the job details, drill down on a job from the **My Jobs** page in the Data Integration service, the **All Jobs** and **Running Jobs** pages in the Monitor service, or from the **Mass Ingestion** page in Operational Insights service.

For an initial load job, the Agent Name is displayed in the **Object Detail** pane. For an incremental load job, the Agent Name is displayed in the **Overview** section of the **Task Summary** pane. For combined initial and incremental load jobs, the Agent Name is displayed in both the **Object Detail** pane and the **Overview** section.

This enhancement was introduced to help identify the Secure Agent for which logs must be collected for troubleshooting when the job runs on a Secure Agent group and one of the tasks fails. Previously, if a task failed and the job was running in a Secure Agent group, you had to manually run the log collector on all of the Secure Agents in the Secure Agent group.

If an incremental load job fails over from one Secure Agent to another, the latest Secure Agent assigned to the job is displayed. If you need to collect logs for a previously assigned Secure Agent, download the complete log from each of the Secure Agents in the Secure Agent group.

## Mass Ingestion Databases

The May 2024 release of Database Ingestion and Replication includes the following new features and enhancements:

### Data validation for initial load jobs that have an Oracle or SQL Server source and a Snowflake target

**Note:** The data validation feature availability is controlled by means of an organization-level feature flag. If this functionality is not available for your organization but you want to use it, create a request for Informatica Global Customer Support.

For database ingestion initial load jobs that have an Oracle or SQL Server source and a Snowflake target, you can now use a new data validation mechanism to compare the target data with the corresponding source object. Database Ingestion and Replication checks data consistency to ensure that all data has been correctly replicated to the appropriate target columns without any data conversion or truncation.

The data validation option is available only for tasks that have the status of **Completed**.

You can run data validation in the **Object Detail** pane of monitoring job-specific details. To access the **Object Detail** pane, drill down on a job from the **My Jobs** page in the Data Integration service, the **All Jobs** and **Running Jobs** pages in the Monitor service, or from the **Mass Ingestion** page in Operational Insights service.

When you run data validation in the Database Ingestion and Replication service, you will be charged based on the CPU consumption on the system where the Data Validation service runs.

### Enhancements to Db2 for z/OS stored procedure processing

For database ingestion incremental load and combined initial and incremental load jobs that have Db2 for z/OS sources, the following enhancements have been made to Db2 for z/OS stored procedure capture processing:

- You can configure the stored procedure to skip empty URs when collecting log data from the z/OS system to help improve performance. Use the `pxw.cdcreader.ZOS.strproc.skipEmptyUR` custom property.
- During change data capture, Mass Ingestion Databases can now pass up to 16 KB of source table filtering information to the Db2 for z/OS stored procedure. Previously, the maximum was 4 KB. If the amount of filtered data exceeds 16 KB, Mass Ingestion Databases displays an error message.



# Changed behavior

The May 2024 release of Data Ingestion and Replication includes the following changed behaviors.

## Monitoring Mass Ingestion jobs in Data Integration and Monitor services

The Data Integration and Monitor services include the following changes for monitoring Mass Ingestion jobs:

- The **My Jobs** page in Data Integration and the **All Jobs** page in Monitor display all of your jobs or all jobs, respectively, but the retention time for undeployed jobs is based on the Job Log Service (JLS) purge policy.  
Both pages display only the new jobs created from the current release. The undeployed jobs from the previous release are excluded.  
Previously, the **My Jobs** page in Data Ingestion and Replication used to display all jobs without purging. Therefore, the number of jobs listed in the Data Ingestion and Replication service could be greater than in the Data Integration service.
- When you deploy an application ingestion task or database ingestion task, the newest job doesn't automatically appear at the top of the jobs list on the **My Jobs** page in Data Integration and the **All Jobs** page in the Monitor. You can sort jobs by using the **Start Time** property or use the find and filter features to search for a specific job.  
Previously, you could sort jobs from newest to oldest on the **My Jobs** page in the Data Ingestion and Replication service.
- When you click the instance name of a file ingestion job on the **My Jobs** page in Data Integration, the job details page appears. When you click the task name from the job details page, it opens in the edit mode. Previously, when you clicked the task name from the job details page, the details opened only in the view mode for file ingestion tasks in Mass Ingestion.
- When you undeploy a streaming ingestion and replication job, the history of the undeployed tasks appears on both the **My Jobs** page in Data Integration and the **All Jobs** page in Monitor.  
Previously, only the latest tasks were listed.

## The generated name of the job instance contains a hyphen (-) instead of an underscore (\_) before the job instance number

The generated name of the job instance now contains a hyphen (-) instead of an underscore (\_) between the task name and job instance number on the **My Jobs** page, on the **Running Jobs** and **All Jobs** pages in Monitor, and in the **Recent jobs** panel of the unified Home page:

```
<taskname>-<job_instance_number>
```

In the Cloud Data Ingestion and Replication Command Line Interface (CLI) and in Operational Insights, the format of the job instance name continues to use the underscore (\_) before the job instance number:

```
<taskname>_<job_instance_number>
```

In the CLI, you must use the `<taskname>_<job_instance_number>` naming convention.

## Default age-out period for incremental load processing has increased

Database ingestion incremental load and combined initial and incremental load jobs stop tracking transactions for restart purposes when those transactions have no DML for tables of interest after an age-out period. The age-out period has been increased from 60 minutes to 24 hours.

The restart point where Data Ingestion and Replication starts reading change data might include recapturing empty transactions that are up to 24 hours old.

# INDEX

## B

behavior changes

- Data Ingestion and Replication August 2024 [9](#)
- Data Ingestion and Replication July 2024 [12](#)
- Data Ingestion and Replication October 2024 [6](#)
- Mass Ingestion May 2024 [17](#)

## D

Data Ingestion and Replication

- changed behavior in August 2024 [9](#)
- changed behavior in July 2024 [12](#)
- changed behavior in October 2024 [6](#)
- new features in August 2024 [7](#)
- new features in October 2024 [4](#), [11](#)

## M

Mass Ingestion

- changed behavior in May 2024 [17](#)
- new features in May 2024 [14](#)

## N

new features

- Data Ingestion and Replication August 2024 [7](#)
- Data Ingestion and Replication October 2024 [4](#), [11](#)
- Mass Ingestion May 2024 [14](#)