



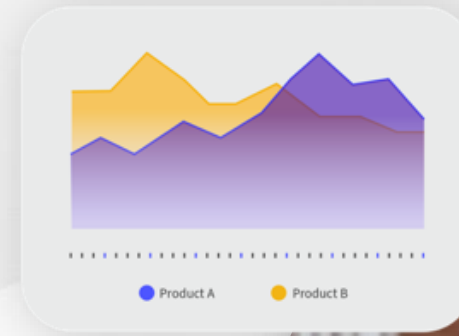
Infoteam – Den store teknikkdagen 2026

Data Fabric

Improve data control with Data Lake Publisher and Data Ledger ++

Per Melander
Industry & Solution Strategy Director

27 January, 2026



CLOUDSUITE SOLUTIONS



INDUSTRY SPECIFIC

Disclaimer

This course provides net-change training content for enhancements delivered in this release. The expectation is that participants have sufficient knowledge of the solution prior to participation in any release training. Release training is not offered for all enhancements. In the event Infor determines, in its sole discretion, not to offer release training for a particular enhancement, there should be sufficient content within the applicable Release Notes and other documentation provided with the release to provide an explanation of the change in the solution regarding such enhancement. Infor does not offer release training for maintenance updates. In addition, the training offered in this course is not intended to be a complete presentation of all problems and issues that may arise in connection with any enhancement. It is also generic and may not be accurate for all. Infor therefore makes no guarantees, or assumes any responsibility for the correctness, sufficiency or completeness of the material presented.

All business entity names, characteristics, products and other data portrayed in this production are fictitious. No identification with actual business entities, names, characteristics, products or other data is intended or should be inferred.

No part of this course may be reproduced or transmitted in any form or by any means electronic or mechanical, without the prior written consent of Infor.

Agenda

1

M3 Publishing data

Data Lake Publisher enhancements

- Table groups and content management
- Initial load or partial reload managed and controlled within DLP.

2

Data Reconciliation

Data Ledger principles for validating consistency in batched and streamed data in Data Lake

- Streamed data
- ION ingestion
- Batch ingested data

3

Data Fabric

Last year's enhancements and news in 2025.10

- Data Catalog
- Compass
- Stream Pipelines
- Undo Purge

4

Source Data Access and Q&A

- Future features in development.
- Questions and Answers

Agenda

1

M3 Publishing data

Data Lake Publisher enhancements

- Table groups and content management
- Initial load or partial reload managed and controlled within DLP.

2

Data Reconciliation

Data Ledger principles for validating consistency in batched and streamed data in Data Lake

- Streamed data
- ION ingestion
- Batch ingested data

3

Data Fabric

Last year's enhancements and news in 2025.10

- Data Catalog
- Compass
- Stream Pipelines
- Undo Purge

4

Source Data Access and Q&A

- Future features in development.
- Questions and Answers

Data Lake Publisher Overview

When you have selected the tables to publish in Data Lake Publisher for your solutions...

And want to do an initial load of 250 tables from M3 to Data Lake...

Or reload all data from 2026-01-12 for all M3 finance tables in one division...

Where is the most logical and intuitive place to find that functionality?



Data Lake Publisher Overview

Enables M3 Business Engine to publish data to Infor Data Fabric.

Data is published through streaming- or batch ingestion. ION ingestion using Dataflows for Data Lake is still supported.

Administration Tools > Data Lake Publisher

Requires role [M3-DLSubscriptionAdmin](#)

To run initial load

Administration Tools > Metadata Publisher

Requires role [M3MDP-User](#)

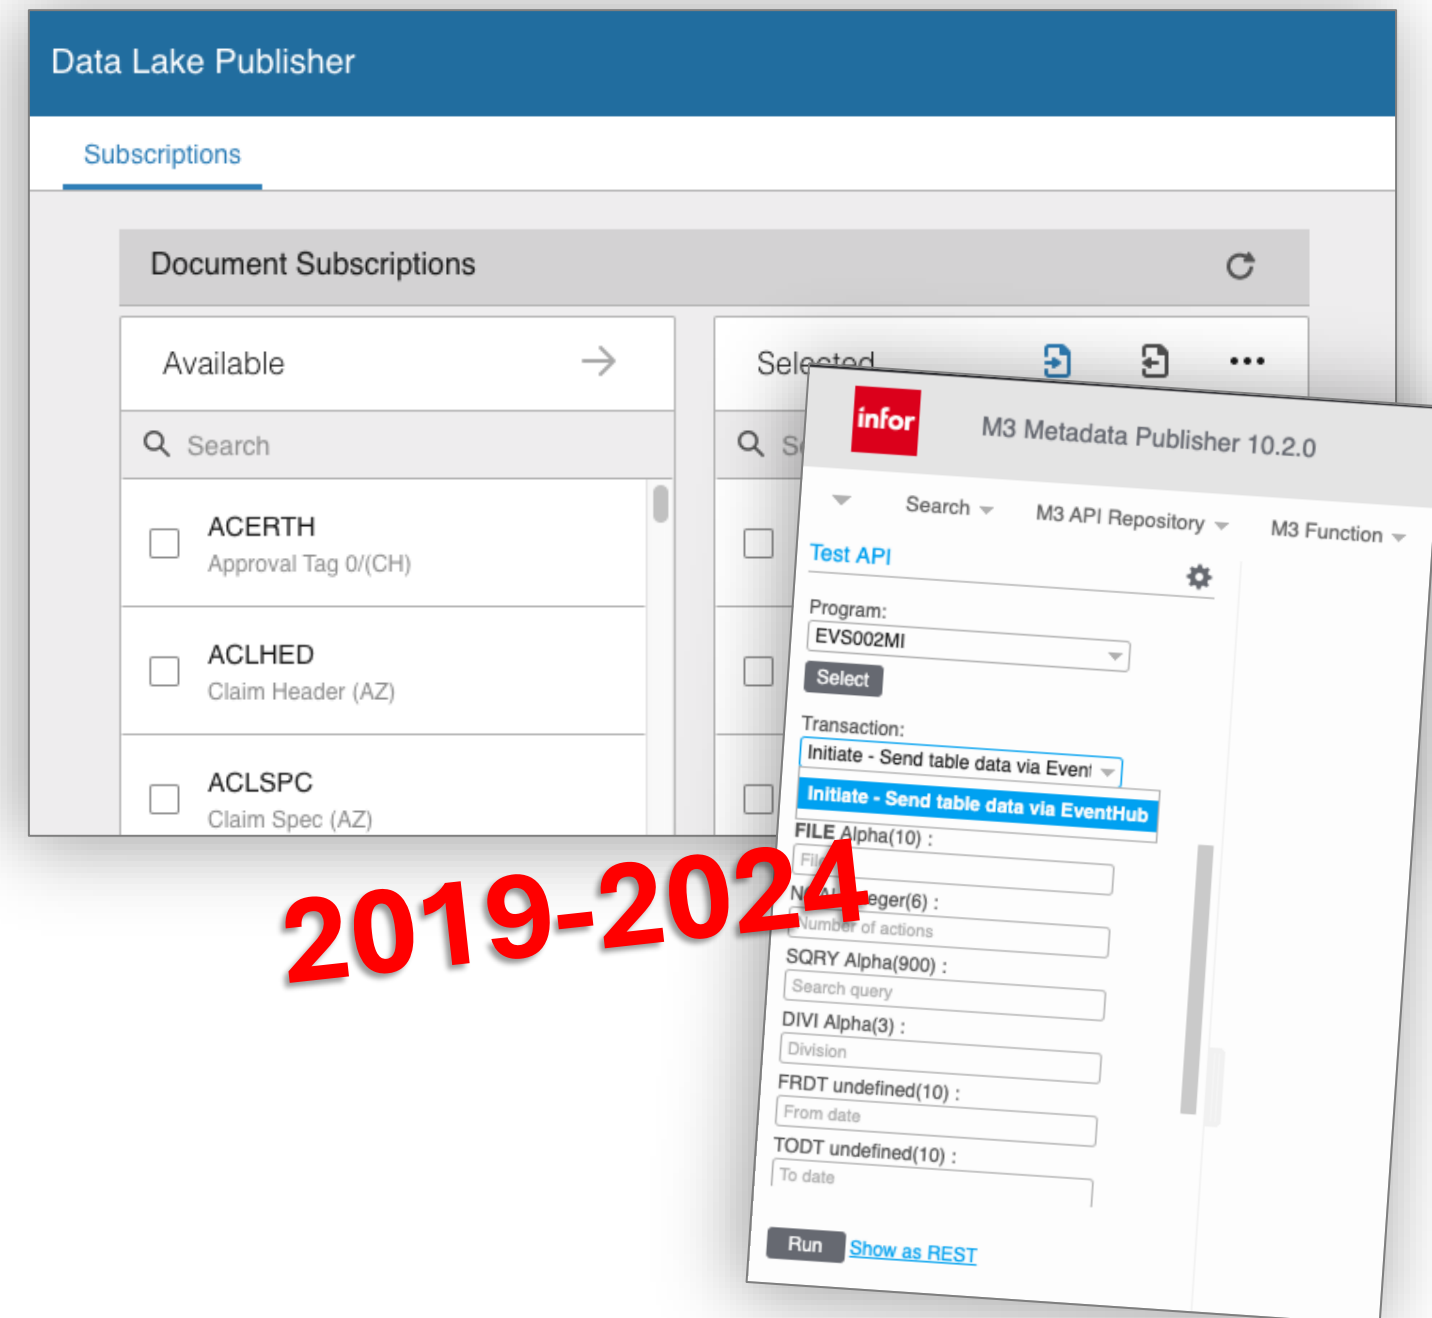


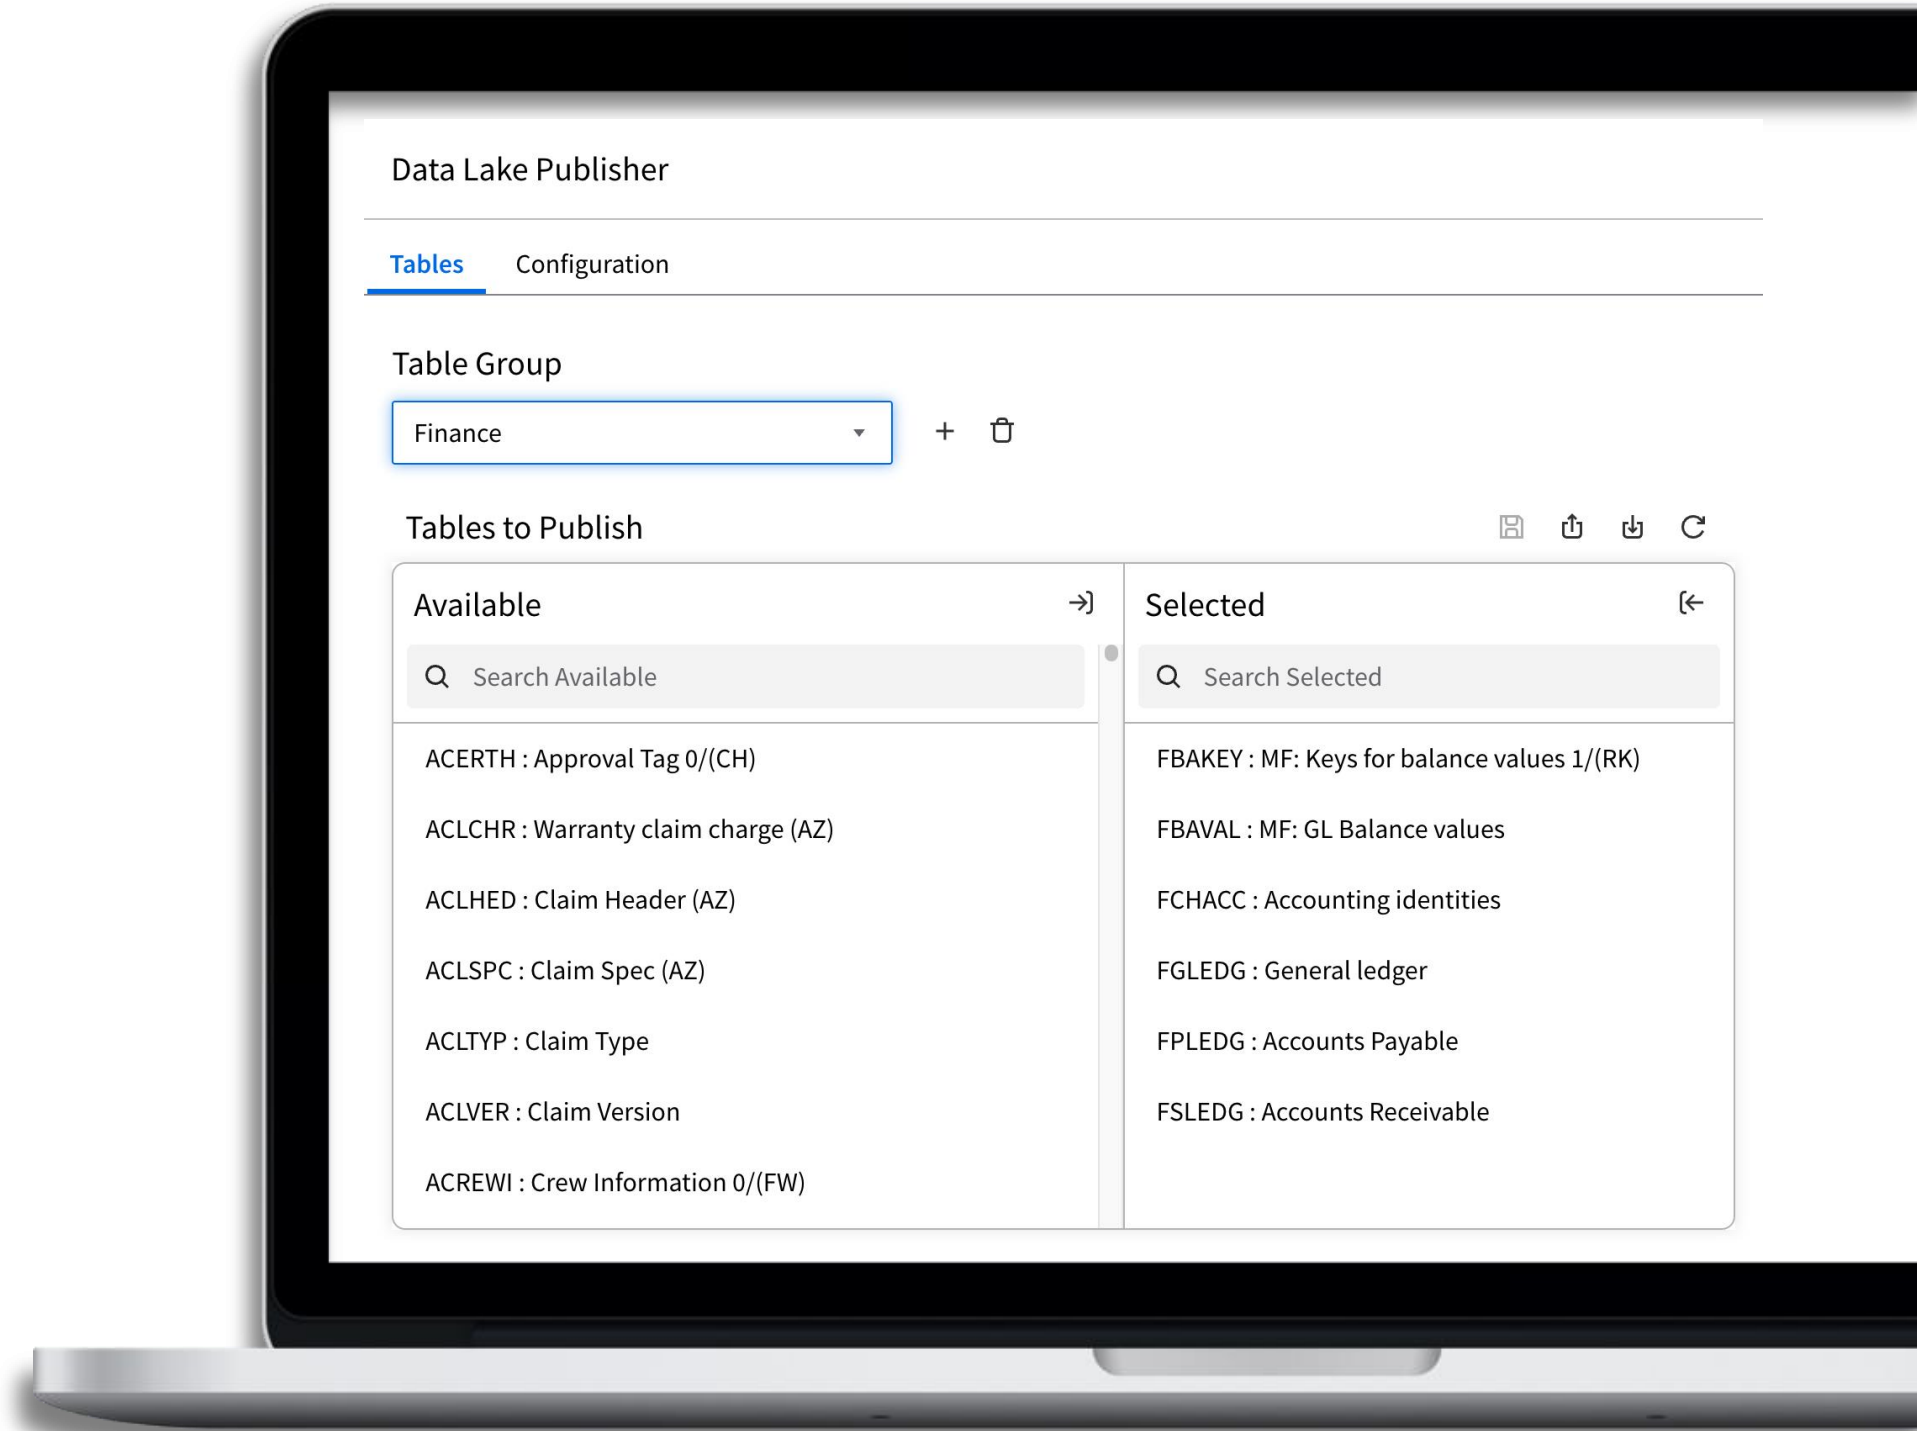
Table Groups

Logical grouping of tables to simplify management and use cases.

- Tables can belong to multiple groups.
- Common use: shared tables for conformed dimensions in analytics.
- Create, view, and delete groups.
- Special groups: (Default) (ungrouped tables) and (All) (all groups in a tree view).

Benefits

- Easier to identify related tables.
- Enables partial updates and exports.
- Supports initial load and partial reload for specific intervals.



M3 Data Lake Publisher - Table Groups

Data Lake Publisher

Tables Configuration

Table Group

(All) + [trash icon]

(Default)

(All)

Finance

Master Data

SPE Package

ACERTH : Approval Tag 0/(CH)

ACLCHR : Warranty claim charge (AZ)

ACLHED : Claim Header (AZ)

ACLSPC : Claim Spec (AZ)

ACLTYP : Claim Type

ACLVER : Claim Version

ACREWI : Crew Information 0/(FW)

ACSALS : Simulation Agreement summary 0/(SS)

ACUACC : Invoice accounting 1/(UT)

ACUACL : Agreement line charges

ACUAEX : Agreement monitoring list

ACUAGA : Maint Agmnt - Handling Chg Item Grp

Selected

Search Selected

- Finance
 - FBAKEY : MF: Keys for balance values 1/(RK)
 - FBAVAL : MF: GL Balance values
 - FCHACC : Accounting identities
 - FGLEDG : General ledger
 - FPLEDG : Accounts Payable
 - FSLEDG : Accounts Receivable
- Master Data
 - CSYCAL : System calendar
 - CSYPAR : System parameters 0/(CP)
 - CSYSTS : States per country
 - CSYTAB : System tables file
 - MITMAH : Style
 - MITMAS : Item Master
 - OCUSMA : Customer
- SPE Package
- default

Data Lake Publisher

Tables Configuration

Table Group

Finance + [trash icon]

Tables to Publish

Available →

Search Available

ACERTH : Approval Tag 0/(CH)

ACLCHR : Warranty claim charge (AZ)

ACLHED : Claim Header (AZ)

ACLSPC : Claim Spec (AZ)

ACLTYP : Claim Type

ACLVER : Claim Version

ACREWI : Crew Information 0/(FW)

Selected ←

Search Selected

FBAKEY : MF: Keys for balance values 1/(RK)

FBAVAL : MF: GL Balance values

FCHACC : Accounting identities

FGLEDG : General ledger

FPLEDG : Accounts Payable

FSLEDG : Accounts Receivable

General Available
April 2025

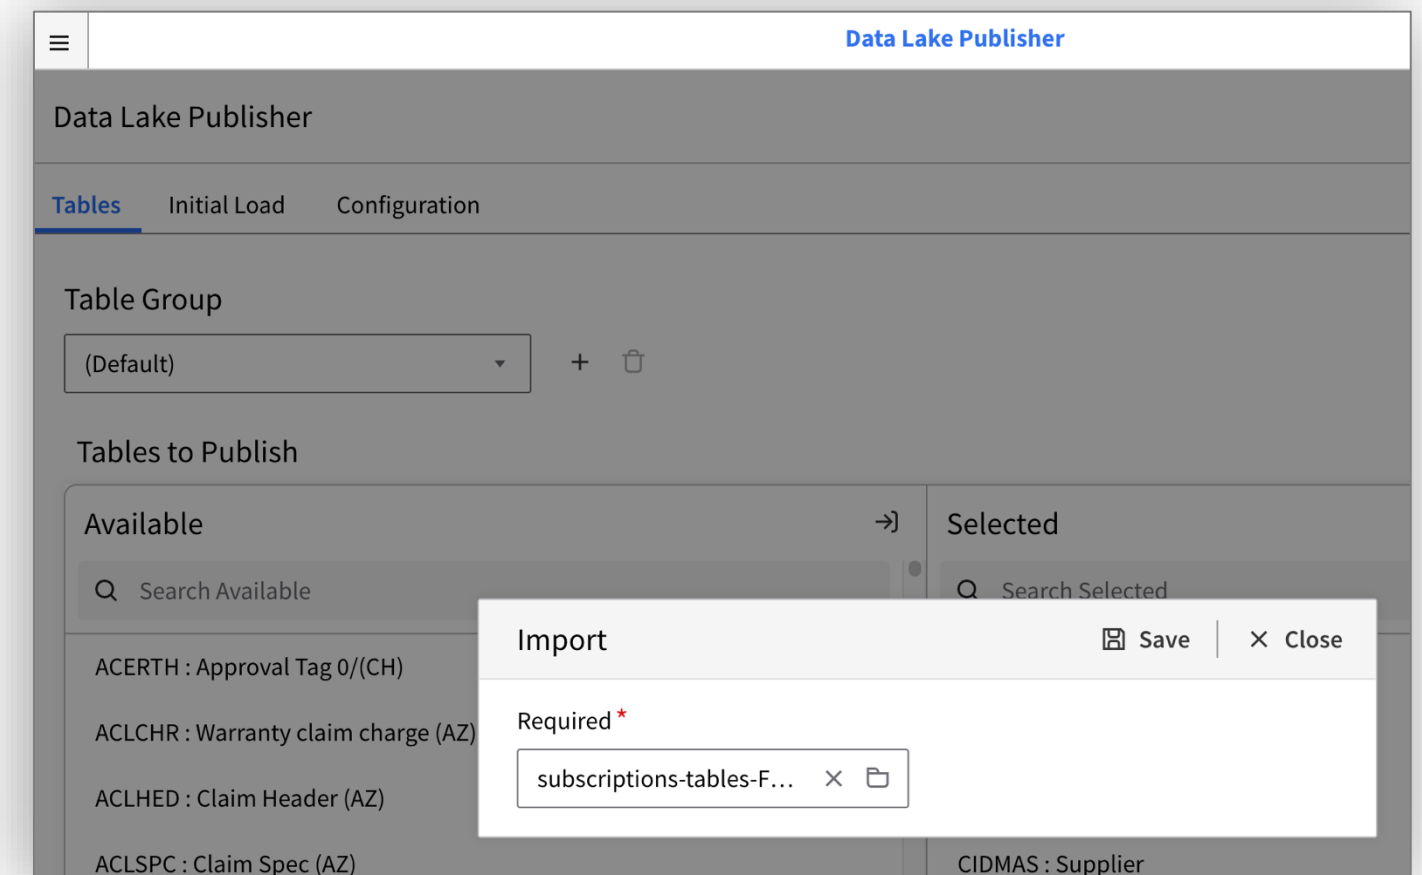
Exporting and Importing Table Selections

Export:

- Export selected tables as a `.zip` file (contains `.json`)
- Can include one or all table groups.

Import:

- Import table selections from a `.zip` file (predefined solution or previous export).
- Automatically creates groups and adds tables as needed.
- Requires saving after import.



Initial or Periodic Loads to the Data Lake

A new user interface allows users to start and monitor initial or periodic data loads to the Data Lake. This removes the need for manual execution and improves visibility and control over the process.

Change Delivery	CloudSuite(s)	Potential Business / Tech. Impact
	All M3 CloudSuites	Medium
Additional Info	Decision maker	Release Training
2025.10	IT Manager	

Problem



- Initial loads required manual API execution (e.g. EVS002MI).
- Job failures occurred when concurrency exceeded 5, with no feedback.
- No user interface to manage or monitor the load process.

Solution



- New tab in Data Lake Publisher for initial load.
- Visual status tracking: queued, running, completed and failed.
- Built-in throttling: enforces a 5-job parallel execution limit automatically.
- Support on-demand incremental reloads (e.g. past month or year).
- Visual status also includes a new tab for job history.

Benefits



- Faster, cleaner process for initial and periodic loads.
- No scripting skills required - fully UI-driven.
- Improved transparency with real-time job tracking.
- Reduces errors by controlling job concurrency.
- Empowers customers to run reloads with minimal effort or downtime.

Select Tables for Initial Loading

Table	Total Row Count	History
- Conformed Dimensions		
ACLCHR : Warranty claim charge (AZ)	0	
CSYCAL : System calendar	66,838	
CSYPER : System periods 0/(CP)	7,911	
CSYTAB : System tables file	23,836	
MITMAS : Item Master	76,420	
OCUSMA : Customer	151,166	
+ default		

Initial Loading in Progress

Submitted	Table	Status	Cancel	Company	Division	From Date	To Date
2025-08-13 16:44:54	ACLCHR : Warranty claim charge (AZ)	WAITING		750			
2025-08-13 16:44:54	CSYCAL : System calendar	WAITING		750			
2025-08-13 16:44:54	CSYPER : System periods 0/(CP)	WAITING		750			
2025-08-13 16:44:54	CSYTAB : System tables file	WAITING		750			
2025-08-13 16:44:54	MITMAS : Item Master	WAITING		750			
2025-08-13 16:44:54	OCUSMA : Customer	WAITING		750			

Initial Loading History

End Time	Duration	Table	Row Count	Status	Restart	Company	Division	From Date	To Date
2025-08-13 16:45:05	00:00:02	CSYTAB : System tables file	11,724	Completed		750			
2025-08-13 16:45:05	00:00:02	CSYCAL : System calendar	30,680	Completed		750			
2025-08-13 16:45:15	00:00:10	MITMAS : Item Master	76,252	Completed		750			
2025-08-13 16:45:15	00:00:10	ACLCHR : Warranty claim charge (AZ)	?	Completed		750			
2025-08-13 16:45:25	00:00:20	OCUSMA : Customer	151,140	Completed		750			

Page 24 of 24 | 10 Records per page

Initial or partial load from M3 to Data Lake

A new user interface allows users to start and monitor initial or periodic data loads to the Data Lake. This removes the need for manual execution and improves visibility and control over the process.

Problem

- Initial loads required a manual API call to EVS002MI.
- When starting initial load jobs resulting in more than five (5) concurrently running, the jobs were cancelled without visual feedback.
- No user interface to manage or monitor the load process.

Solution

- New tab in the Data Lake Publisher application to manage initial load.
- Visual status tracking: queued, running, completed and failed jobs.
- Built-in throttling: enforces a 5-job parallel execution limit automatically.
- Support on-demand incremental reloads (e.g. past month or year).

Initial or partial load from M3 to Data Lake

A new user interface allows users to start and monitor initial or periodic data loads to the Data Lake. This removes the need for manual execution and improves visibility and control over the process.

Benefits

- Faster, cleaner process for initial and periodic loads.
- No scripting skills required - fully UI-driven.
- Improved transparency with real-time job tracking.
- Reduces errors by controlling job concurrency.
- Empowers customers to run reloads with minimal effort or downtime.

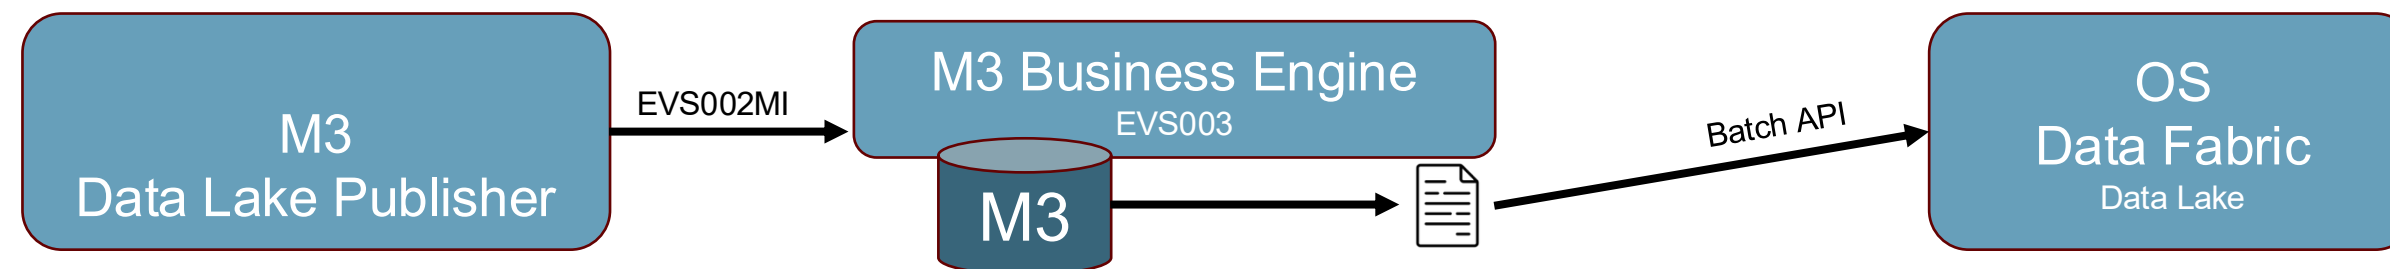
How does it work

Data Lake Publisher (DLP) communicates with the M3 Business Engine to retrieve table information such as table row count and primary keys and columns. This will guide you as a user on where it's possible to filter data by division or using record timestamp incremental extraction.

The M3 Business Engine (M3BE) is called by DLP through EVS002MI to run the publishing job.

EVS002MI starts the M3BE program EVS003 that extracts the table data based on the filter criterias into json files and sends them to Data Fabric batch API.

Data Fabric indexes the data and makes it available to Compass and any active pipeline.



Initial Load using Data Lake Publisher

Tables **Initial Load** Configuration

Select Tables
Progress
History

Select Tables for Initial Loading

<input checked="" type="checkbox"/>	Table	Total Row Count	History
<input checked="" type="checkbox"/>	[A] <input type="text"/>		
<input checked="" type="checkbox"/>	— Conformed Dimensions		
<input checked="" type="checkbox"/>	CFACIL : Facility master	34,753	
<input checked="" type="checkbox"/>	CSYTAB : System tables file	89,536	
<input checked="" type="checkbox"/>	FGINHE : Invoice - head 1/(F4)	68,038	
<input checked="" type="checkbox"/>	MITMAS : Item Master	232,008	
<input type="checkbox"/>	— Table Group Test 1		
<input type="checkbox"/>	— Table Group Test 2		
<input type="checkbox"/>	— Table Group Test 3		
<input type="checkbox"/>	— Table Group Test 4		
<input type="checkbox"/>	— Table Group Test 5		
<input type="checkbox"/>	— Table Group Test 6		
<input type="checkbox"/>	— Table Group Test 8		

Conditions for selected tables

Company* Division

From Date To Date

[Start Initial Load](#)

Initial Loading in Progress

Submitted	Table	Status	Cancel	Company	Division	From Date	To Date
2025-04-01 11:47:27	MITBAL : Item - warehouse	Running		750		2023-12-31	
2025-04-01 11:47:27	MITTRA : Stock transactions	Running		750			
2025-04-01 11:47:27	MMXMAS : Copy of MITMAS 1/(XM)	Running		750		2024-03-31	
2025-04-01 11:47:27	OCUSMA : Customer	Running		750		2024-01-01	2024-04-30
2025-04-01 11:47:27	OINVOH : Invoice Number File	Queued		750		2023-12-31	

Initial Loading History

End Time	Duration	Table	Row Count	Status	Restart	Company	Division	From Date	To Date
2025-04-01 11:47:49	00:00:09	CFACIL : Facility master	25,086	Completed		750		2023-12-31	
2025-04-01 11:47:27	00:00:24	CSYTAB : System tables file	89,536	Completed		750			
2025-04-01 11:47:27	00:00:19	FGINHE : Invoice - head 1/(F4)	50,634	Completed		750	AAA	2024-03-31	2024-03-31
2025-04-01 11:47:27	00:00:03	MITMAS : Item Master		Failed		750		2024-01-01	2024-04-30
2025-04-01 11:47:27	00:00:09	MITBAL : Item - warehouse		Failed		750			
2025-04-01 11:47:27	00:00:17	MMXMAS : Copy of MITMAS 1/(XM)	68,031	Completed		750			
2025-04-01 11:47:27	00:00:33	OCUSMA : Customer	231,055	Completed		750		2023-11-01	2023-11-30
2025-04-01 11:47:27	00:00:09	OINVOH : Invoice Number File		Failed		750			

General Available
October 2025

Agenda

1

M3 Publishing data

Data Lake Publisher enhancements

- Table groups and content management
- Initial load or partial reload managed and controlled within DLP.

2

Data Reconciliation

Data Ledger principles for validating consistency in batched and streamed data in Data Lake

- Streamed data
- ION ingestion
- Batch ingested data

3

Data Fabric

Last year's enhancements and news in 2025.10

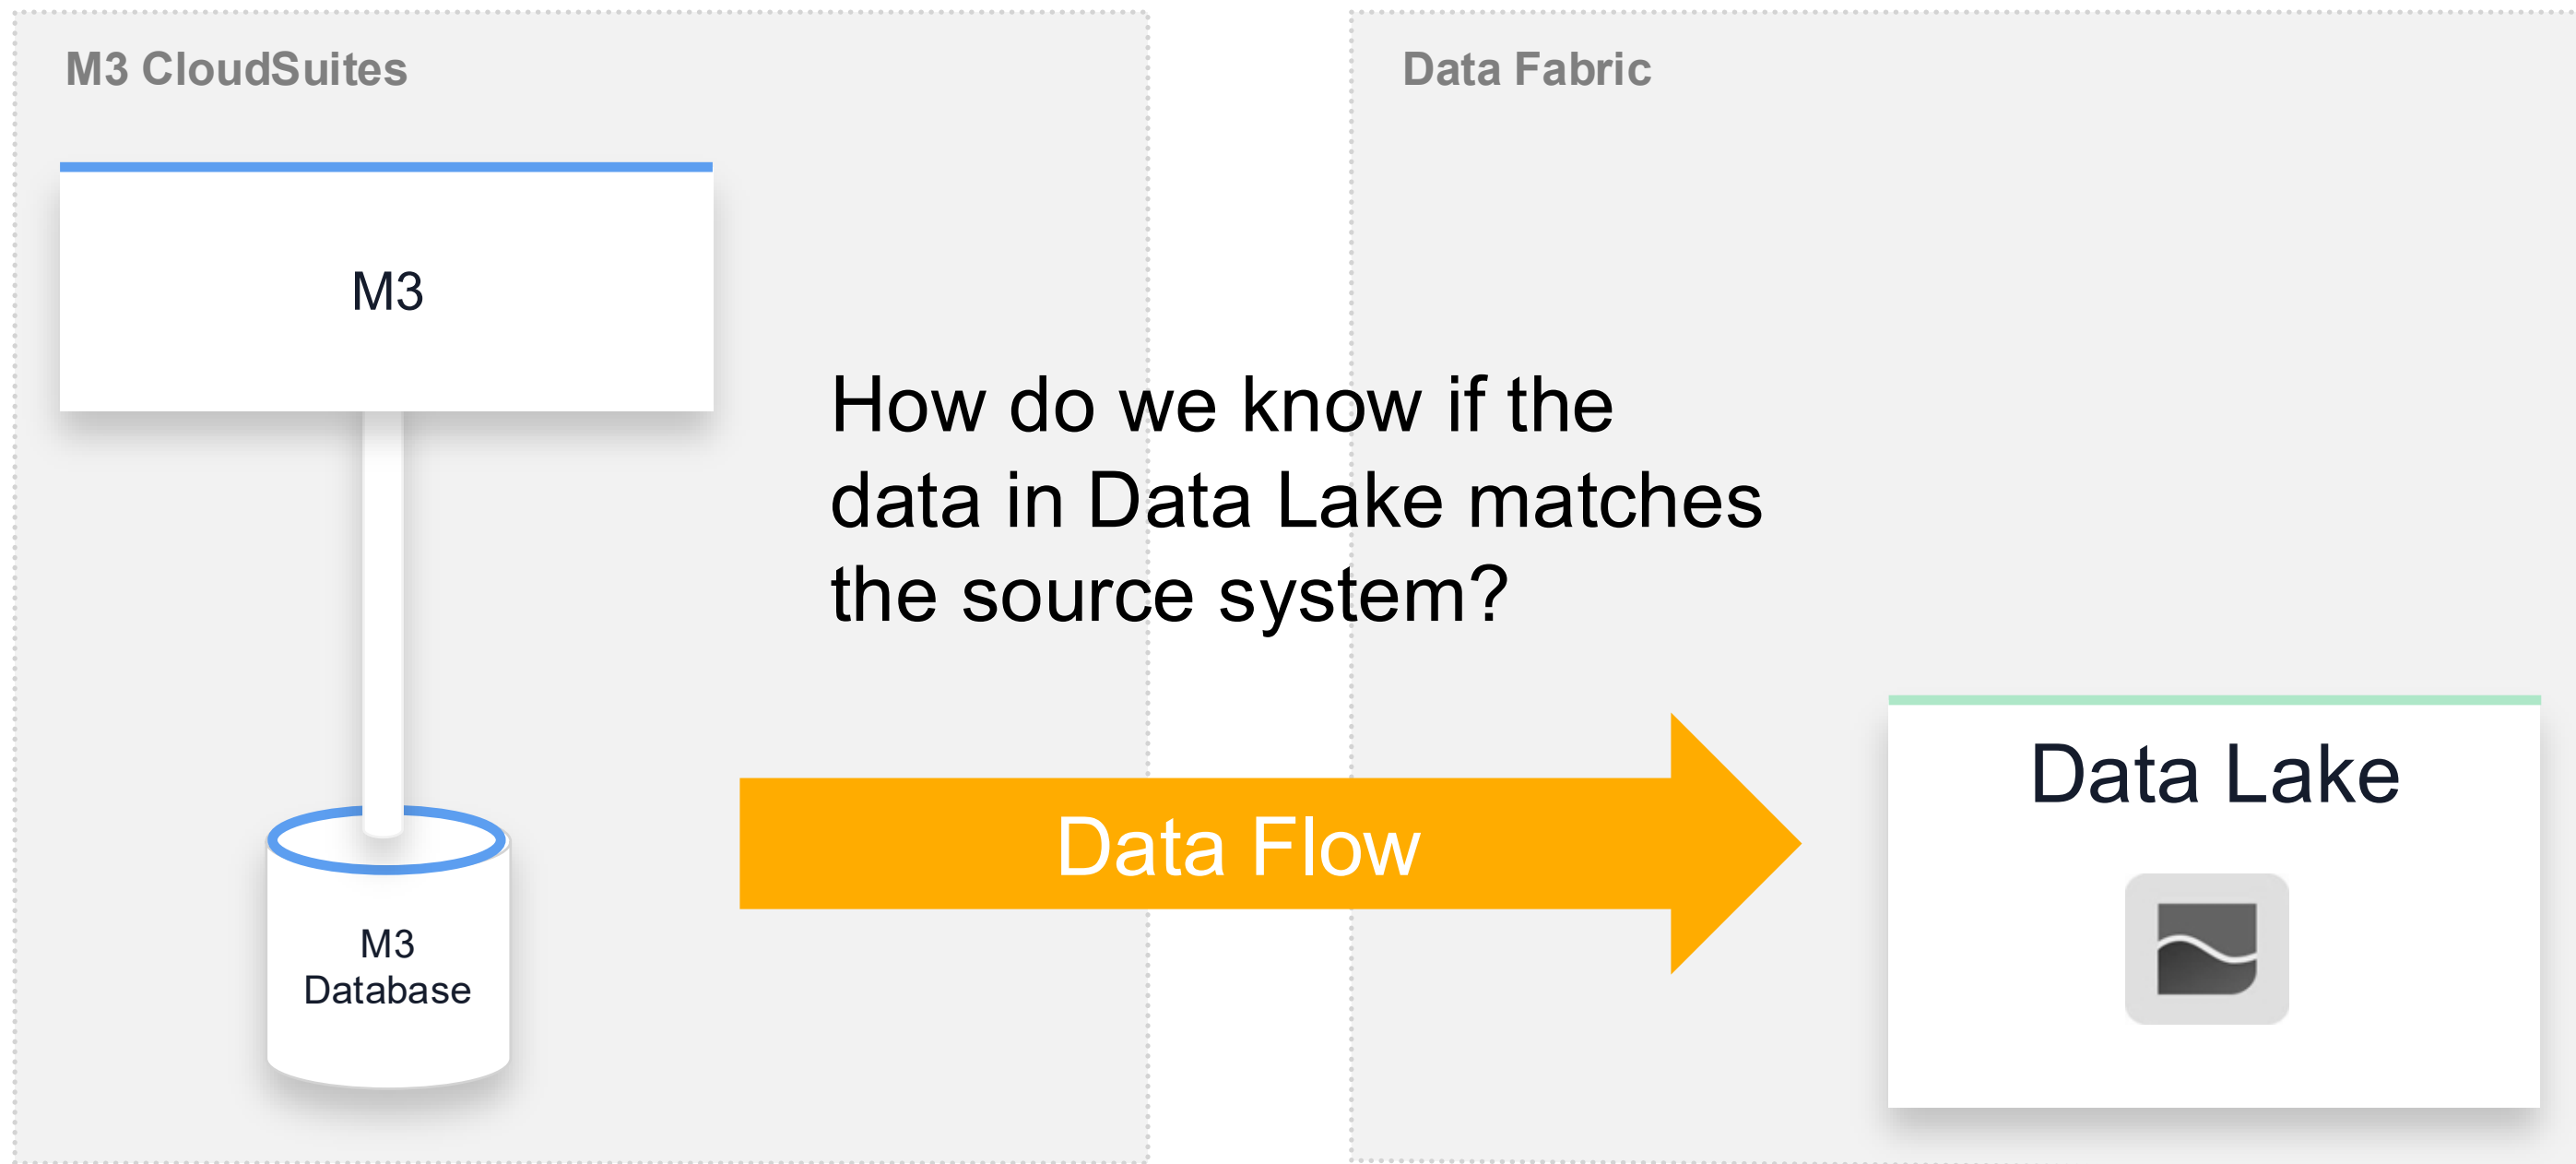
- Data Catalog
- Compass
- Stream Pipelines
- Undo Purge

4

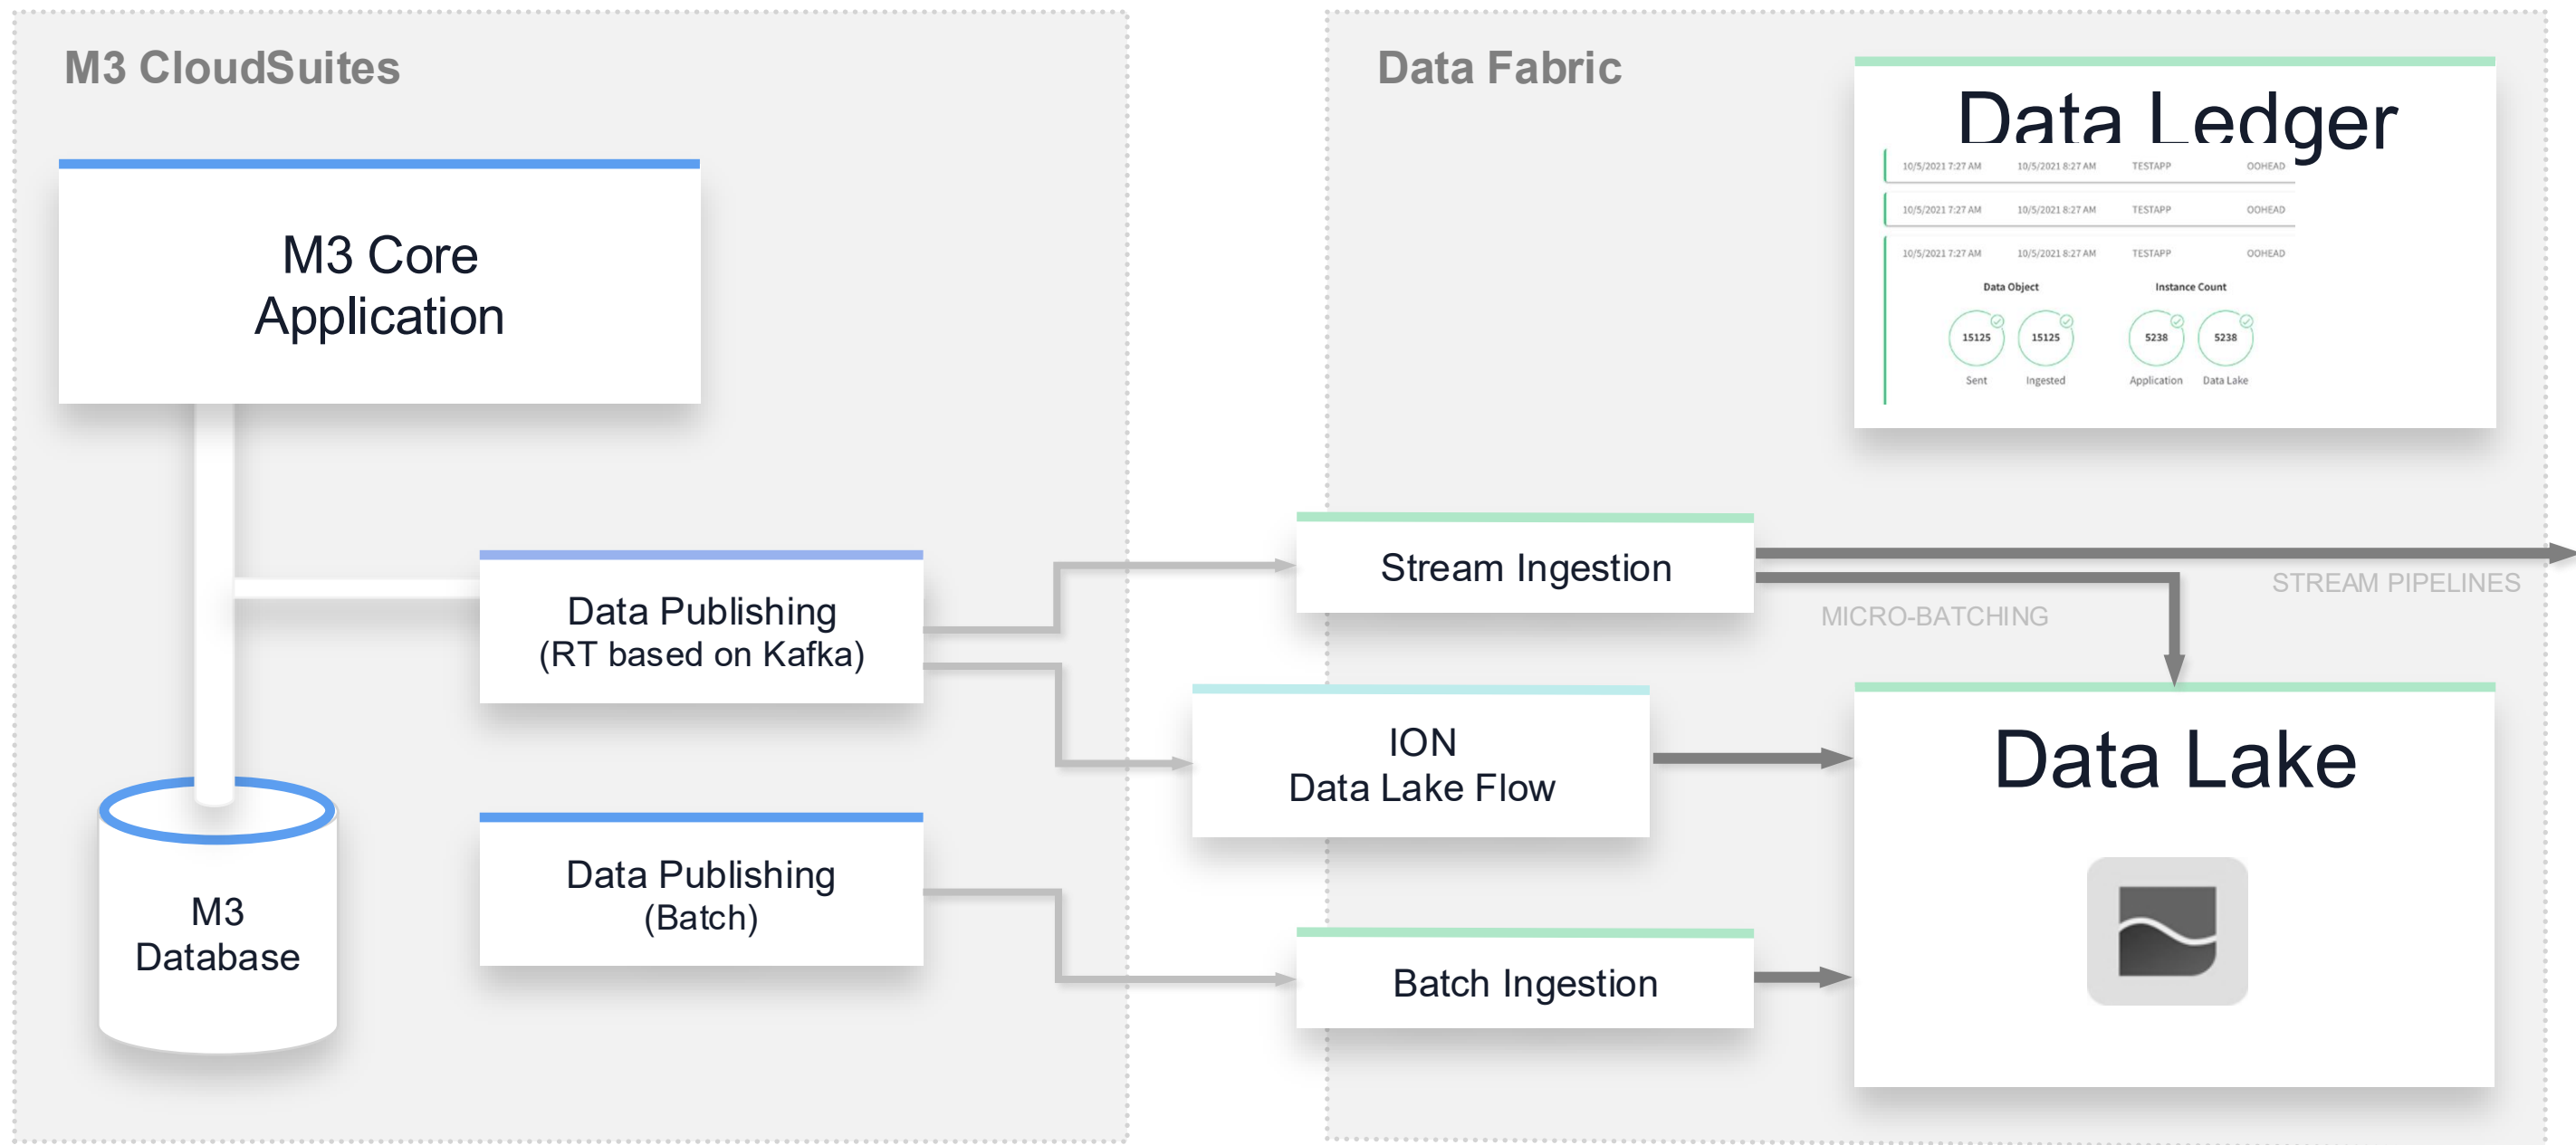
Source Data Access and Q&A

- Future features in development.
- Questions and Answers

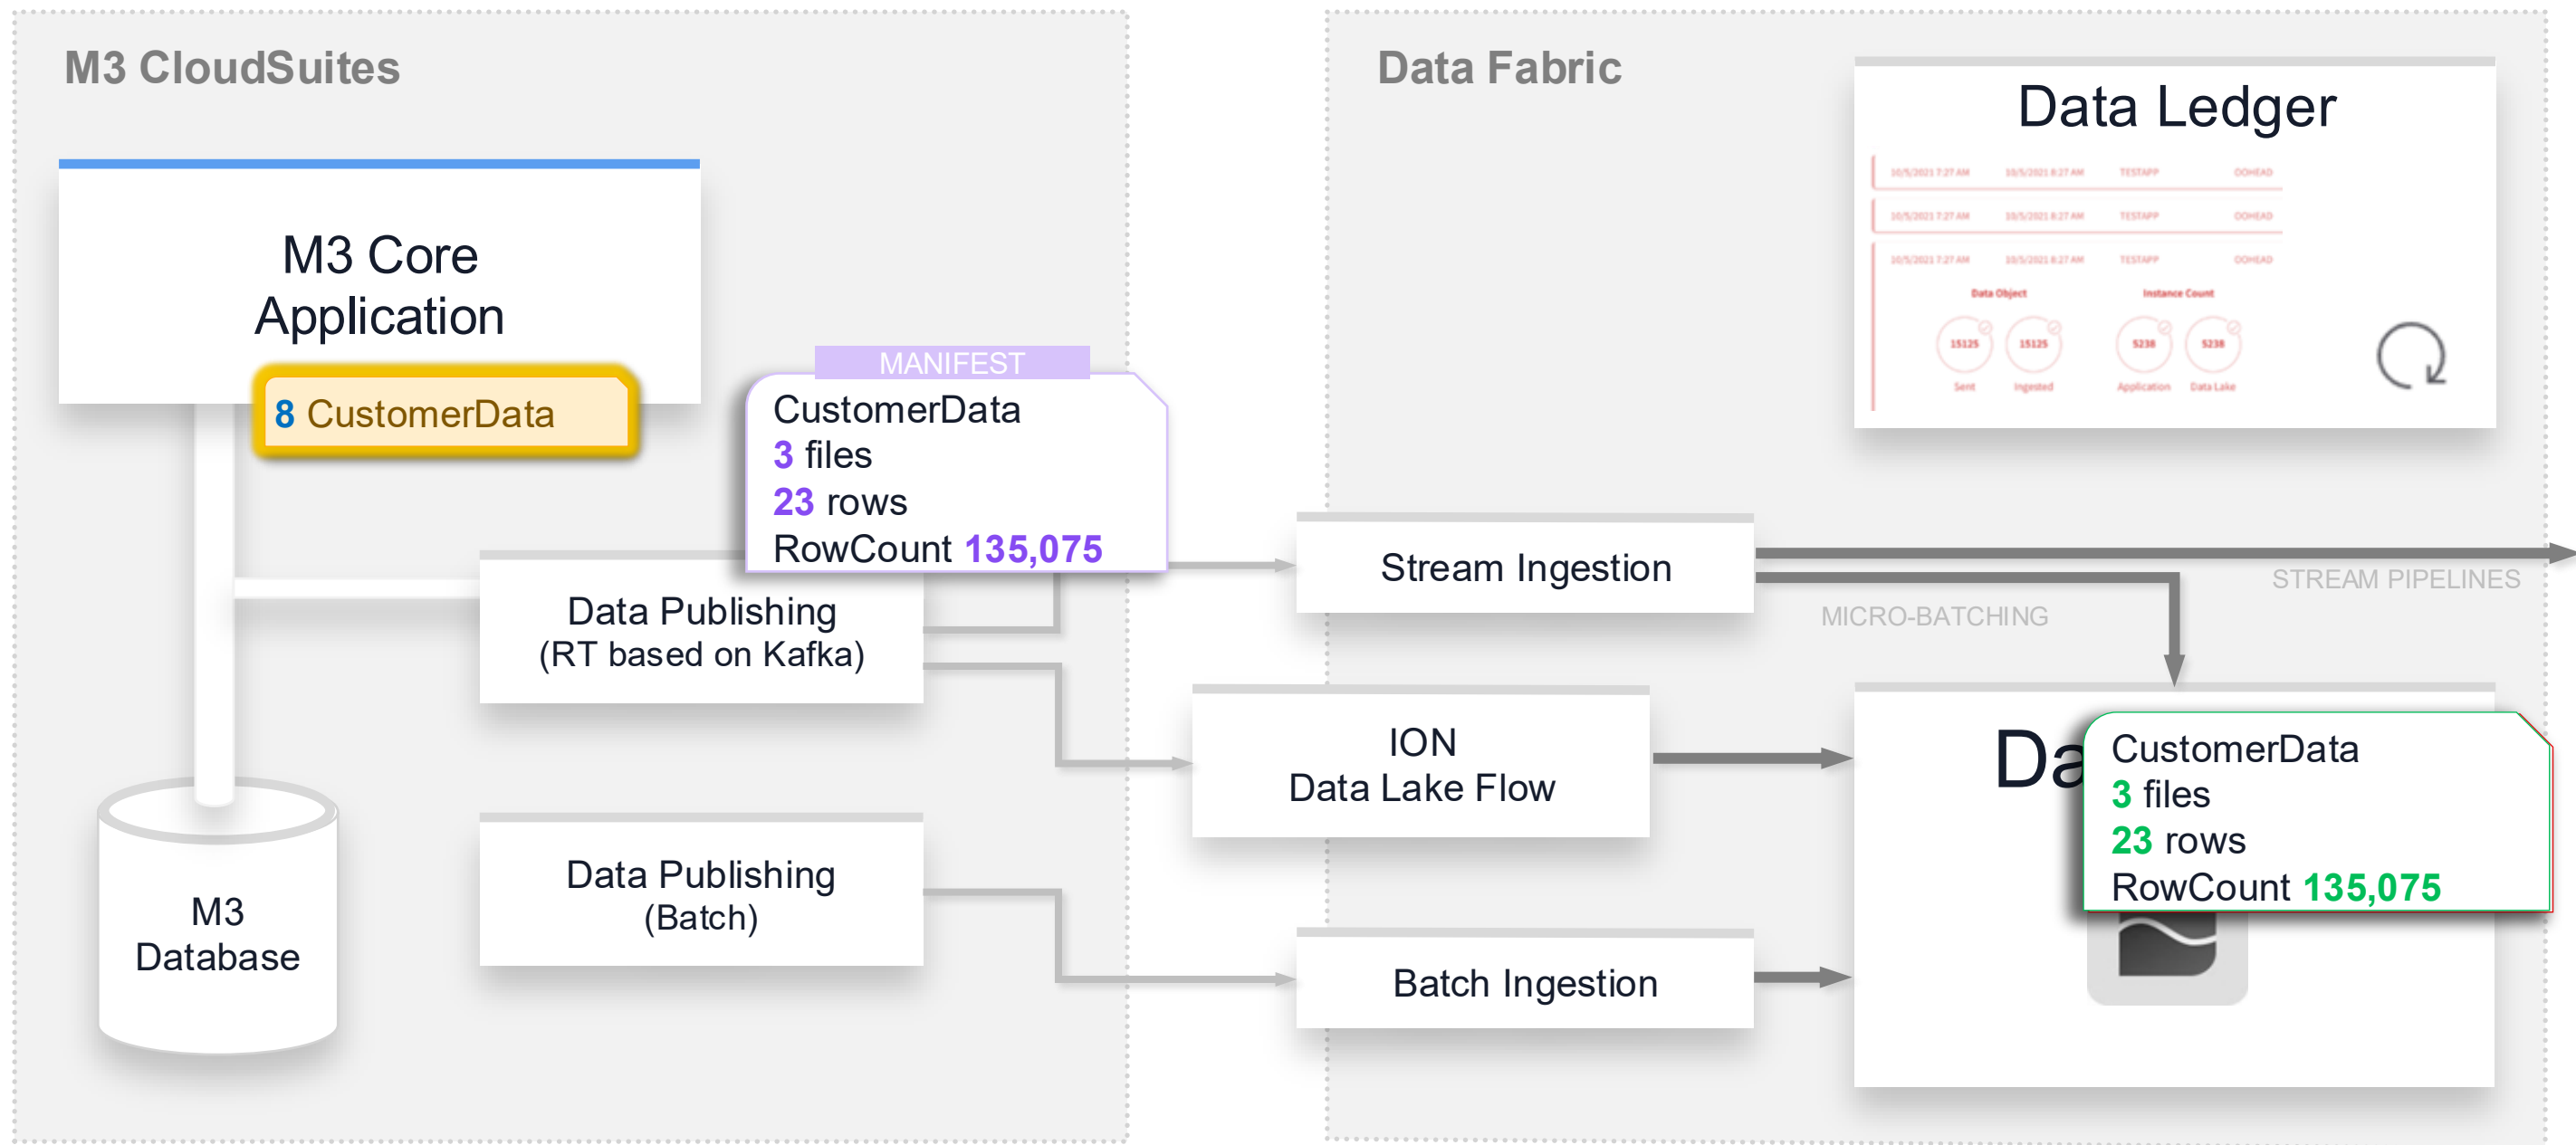
M3 and Data Fabric Consistency



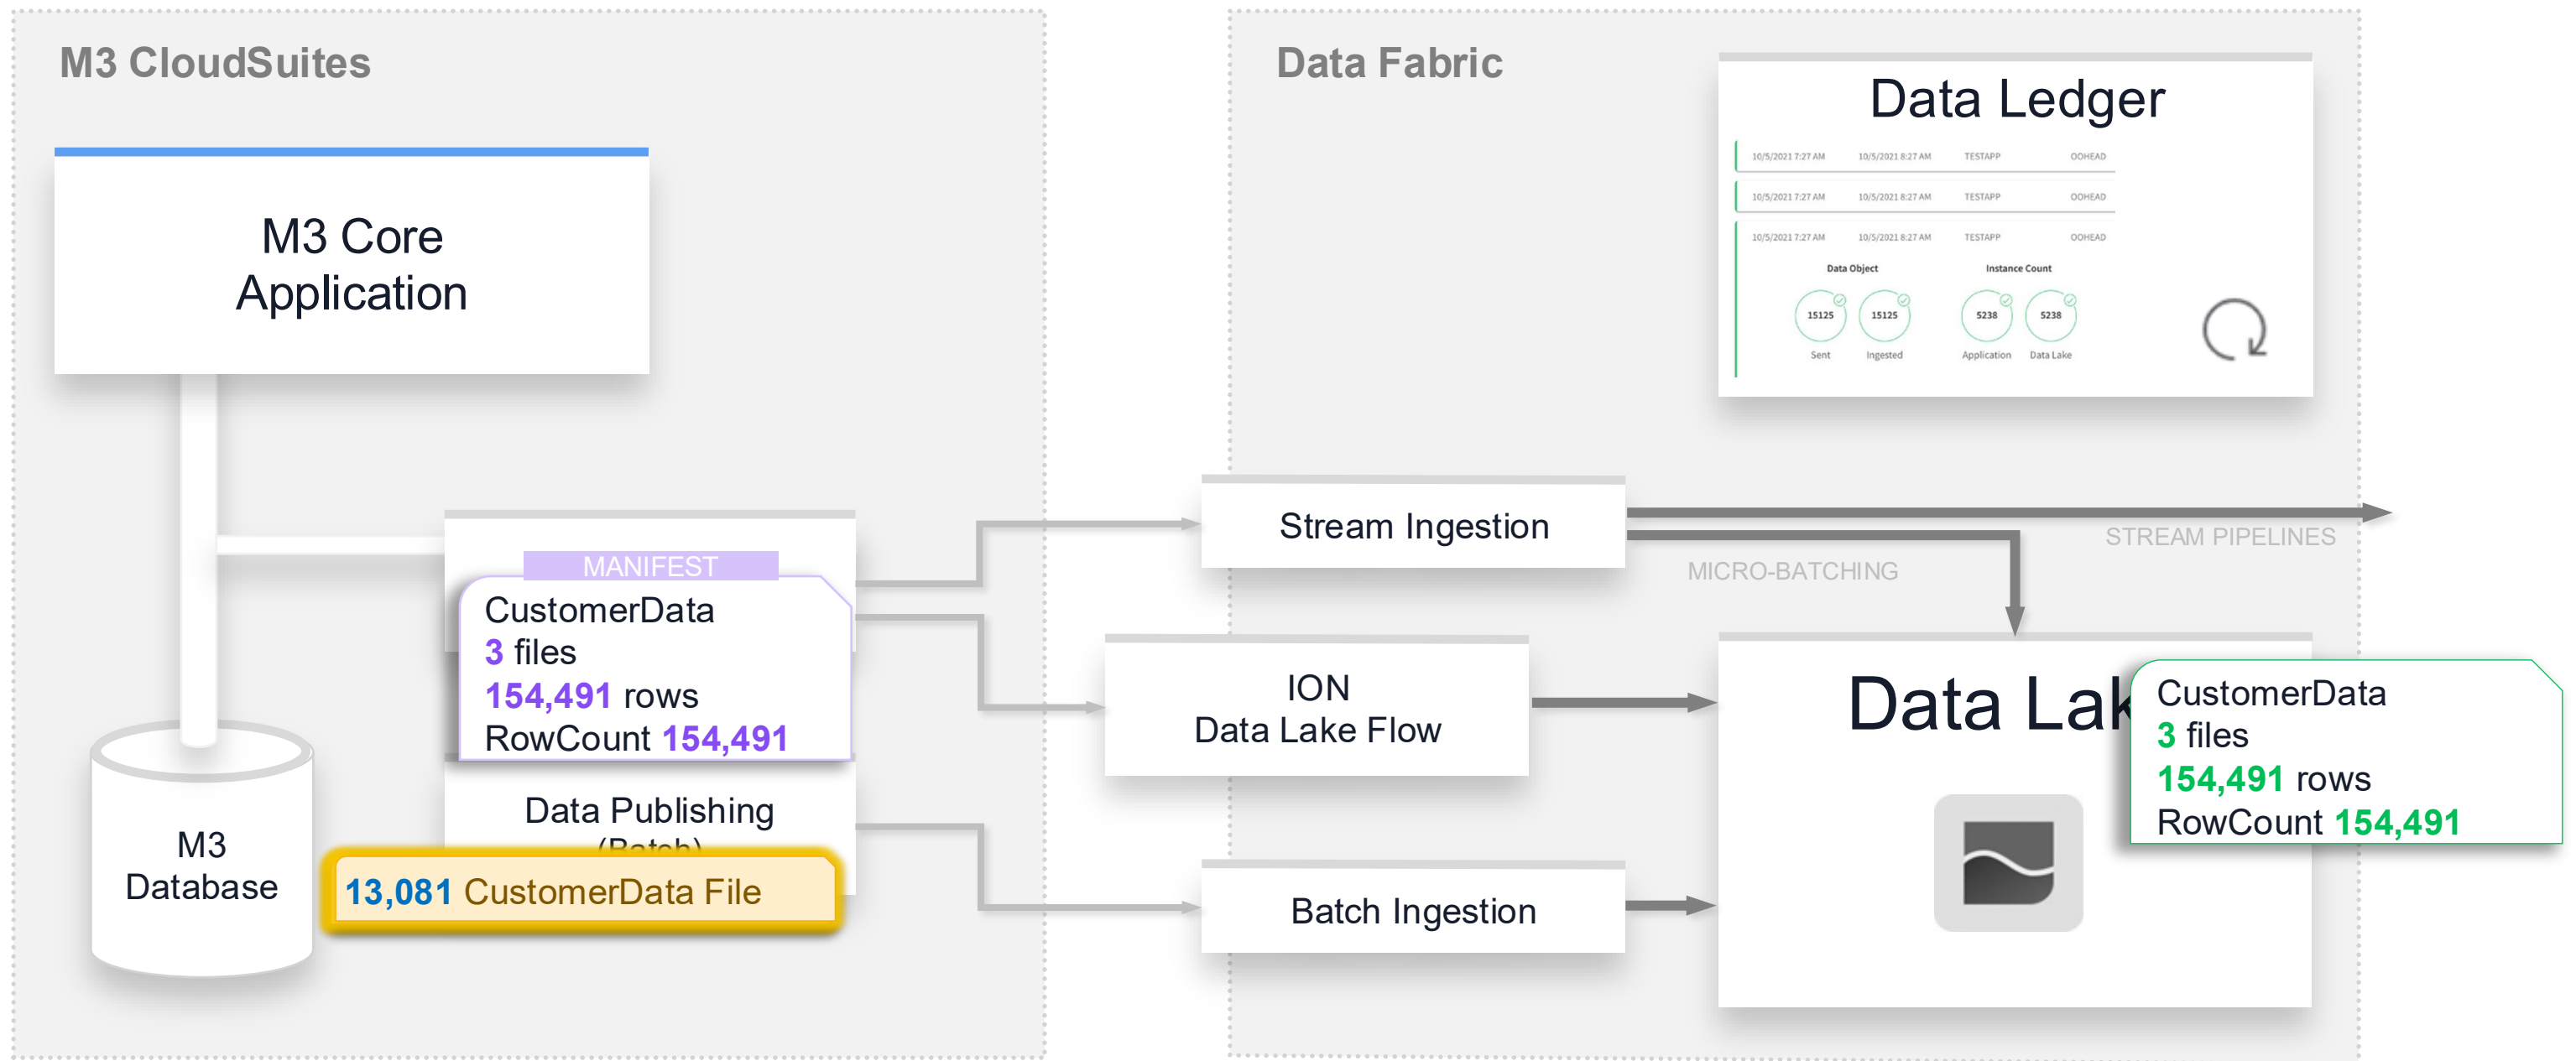
M3 and Data Ledger



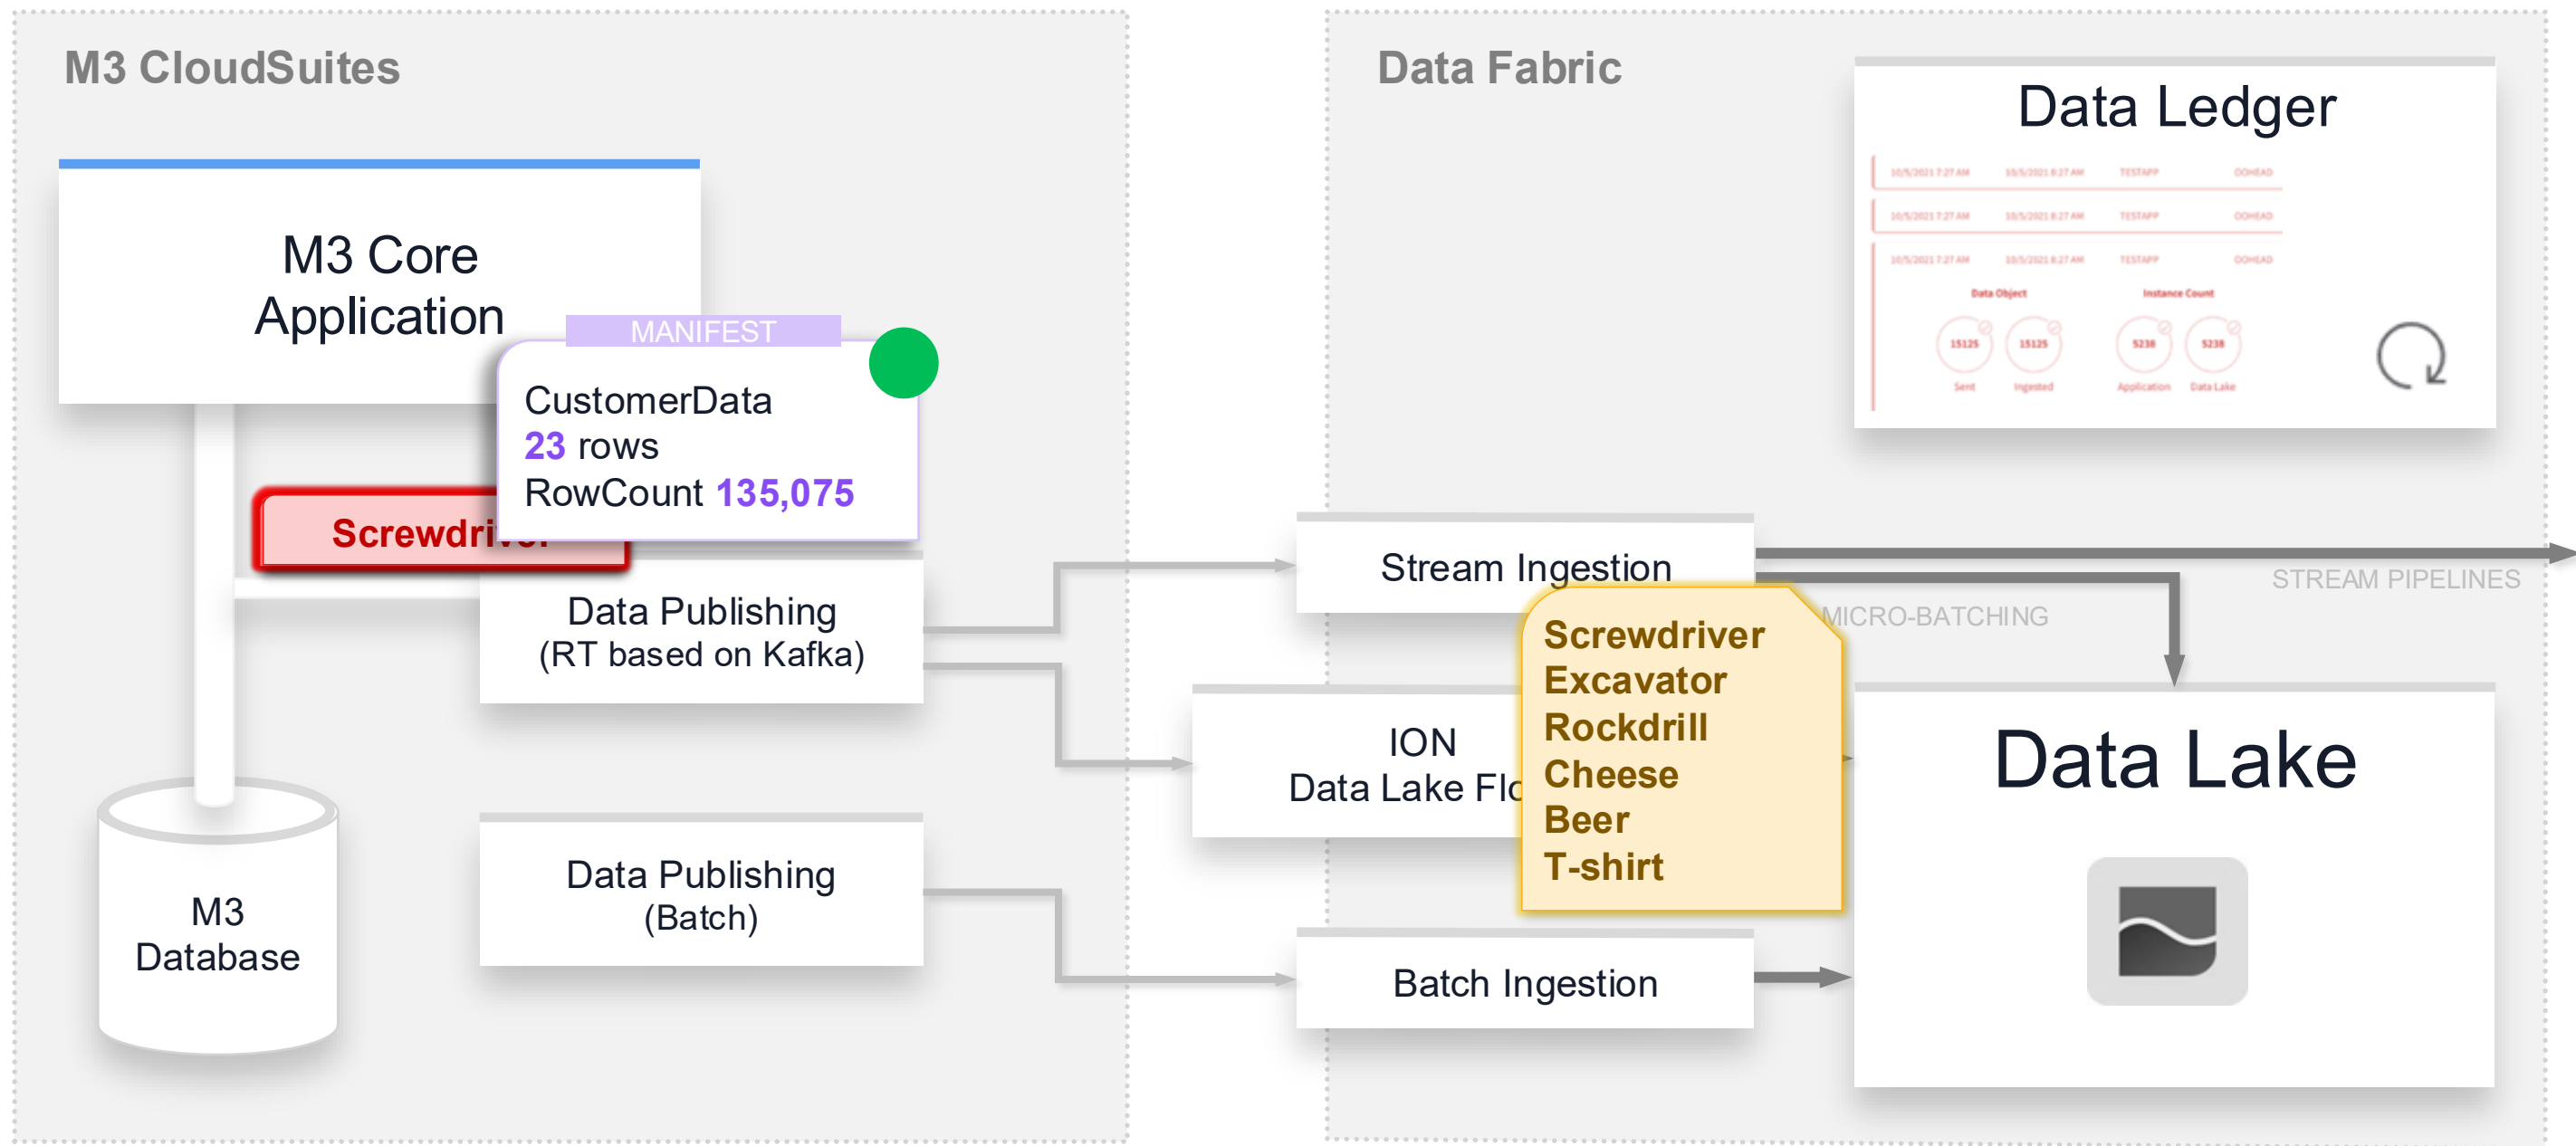
M3 and Data Ledger – Classic ION Ingestion



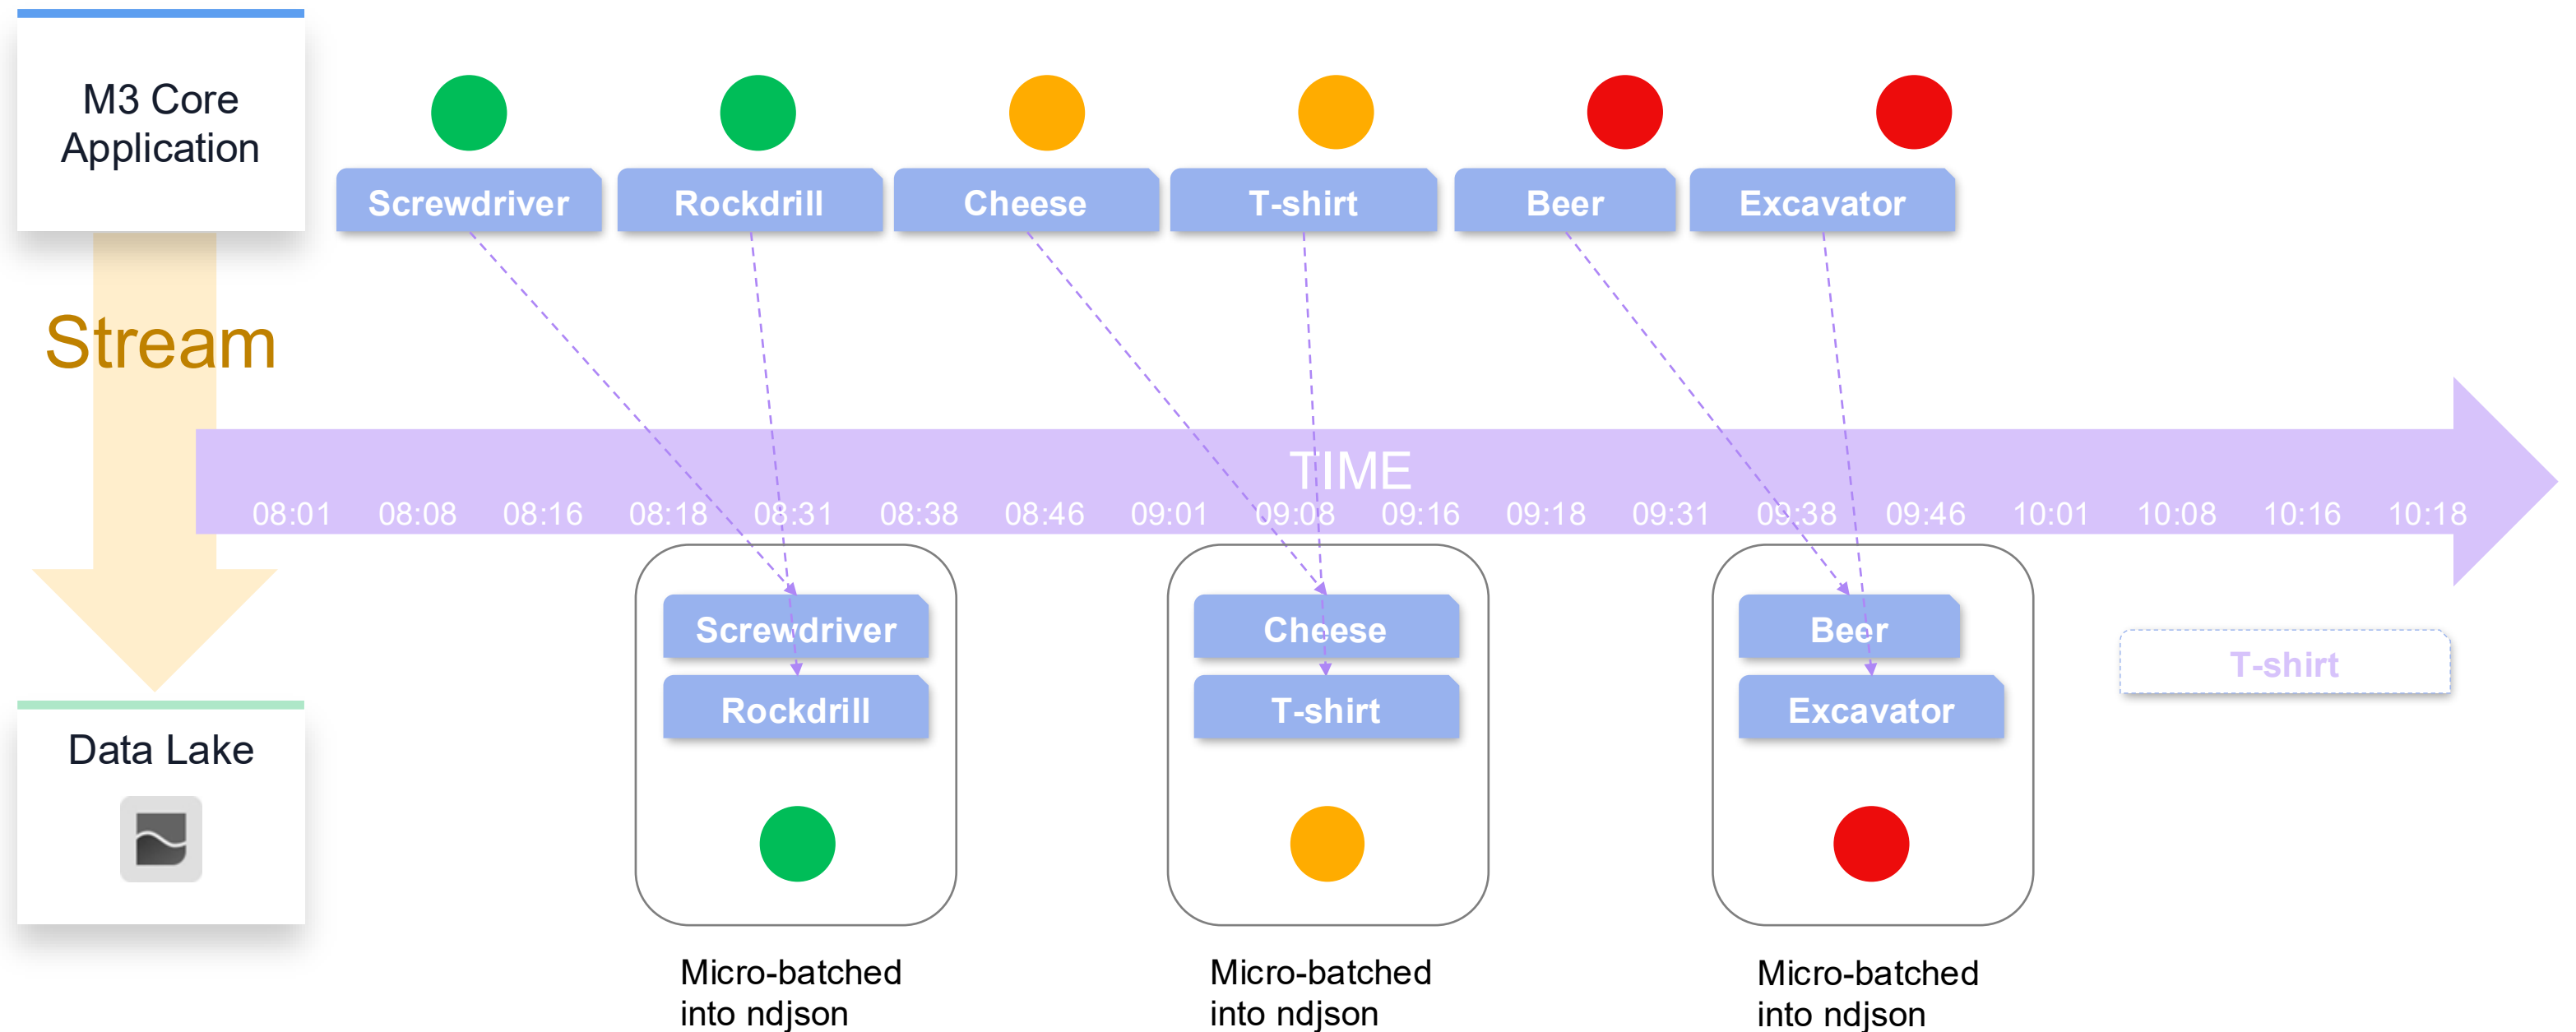
M3 and Data Ledger – Batch Ingestion



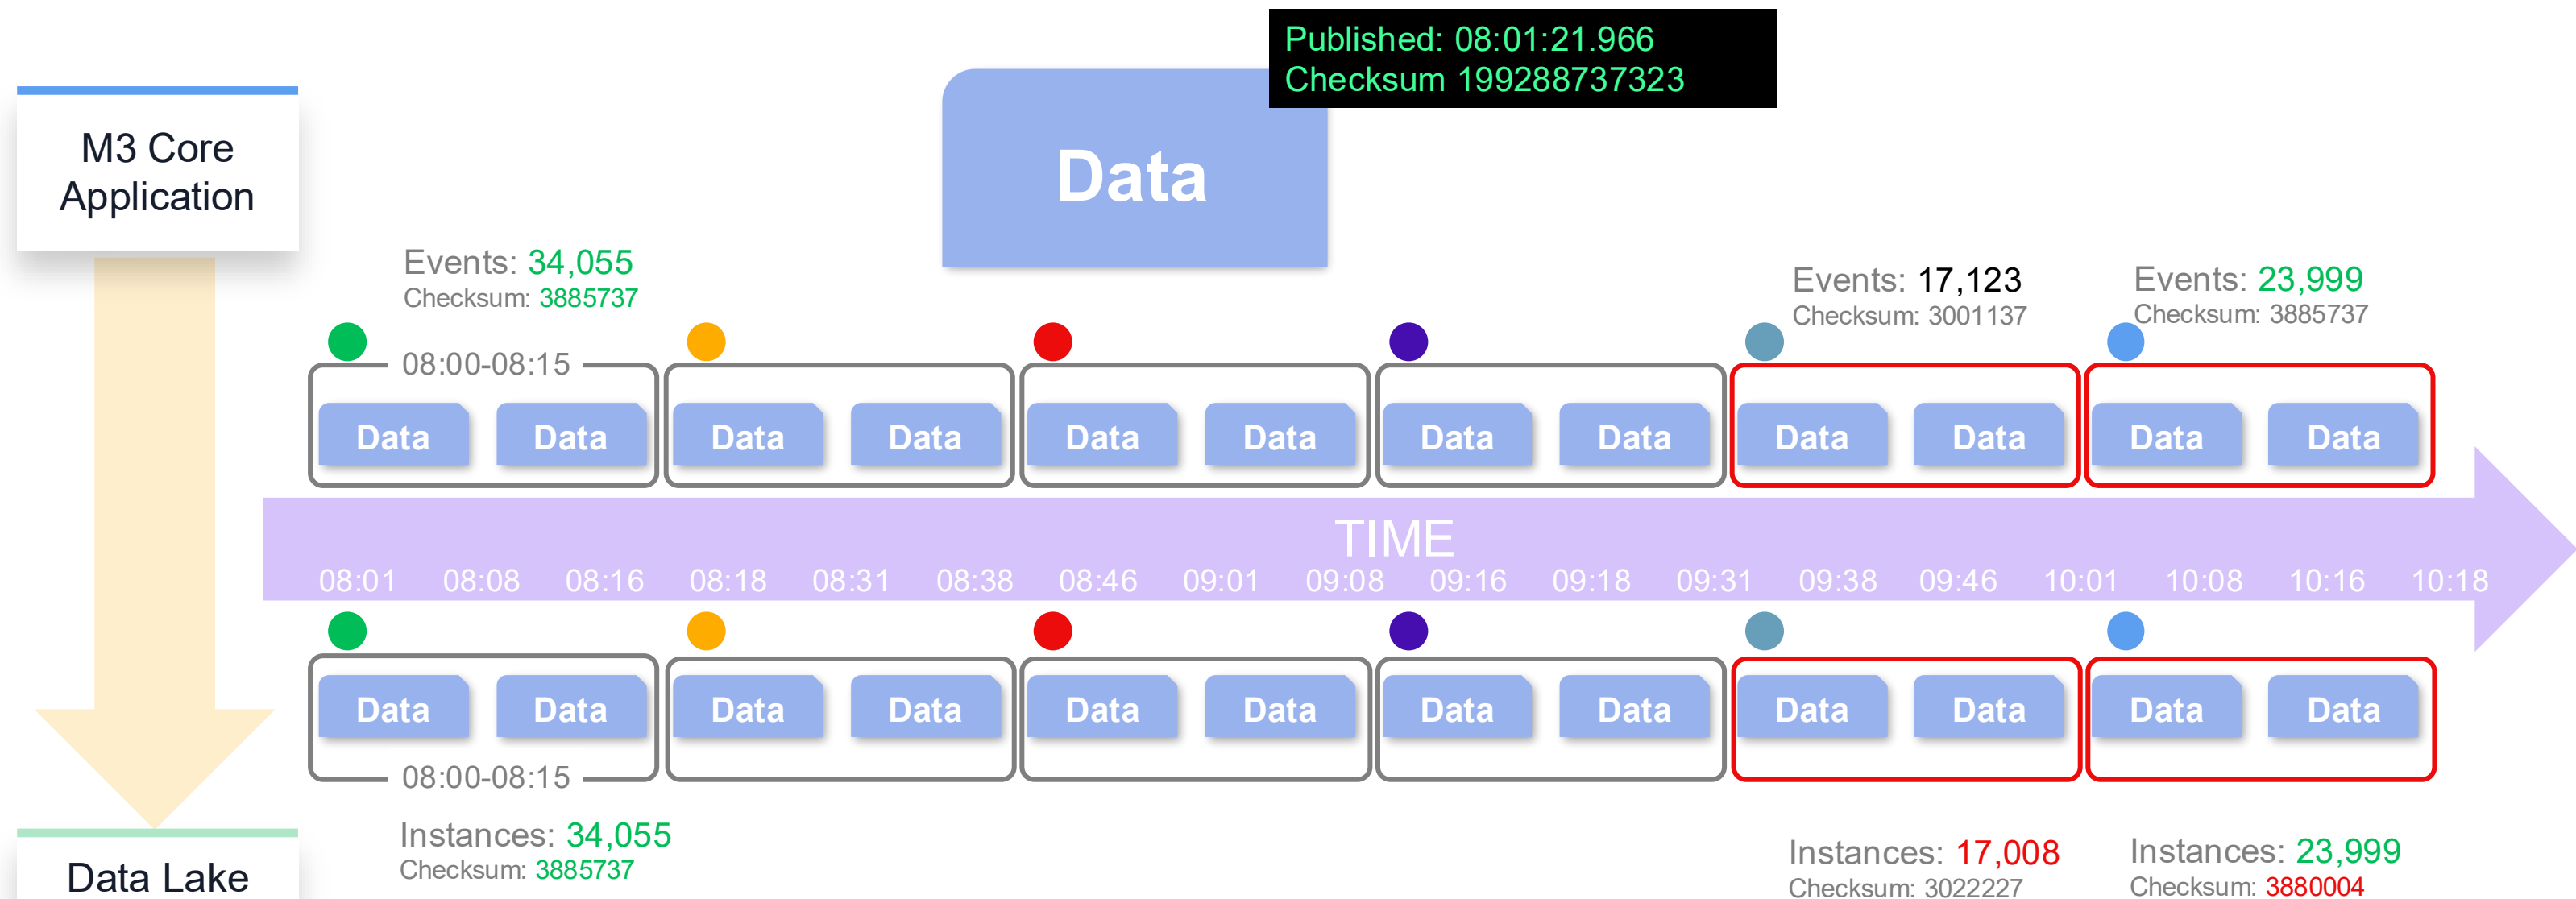
M3 and Data Ledger – Stream Ingestion



How can we compare data in two fast-moving systems



How can we compare data in two fast-moving systems



- Grouping streamed data into virtual batches .
- Within the timeframe of the virtual batch – compare the row counts and the checksums

Data Ledger v2 for Streaming

Redesigned and improved reconciliation feature using checksums

Change Delivery	CloudSuite(s)	Potential Business / Tech. Impact
	All M3 CloudSuites	Medium
Additional Info	Decision maker	Release Training
2025.10	IT Manager	

Problem



- Data is constantly created, transferred and ingested into Data Fabric which makes it difficult to identify and validate the consistency at a specific time
- Admins need to frequently access the Data Ledger page to detect any discrepancies between the system of record and the Data Lake.

Solution



- Streaming statistics are grouped into 15-minute windows for validation on a more detailed level. Checksums on data ensure the correct data being received.
- Administrators can configure Data Ledger notifications through Health Service to promptly identify matches and mismatches in the reconciliation process, ensuring data integrity and trustworthiness.

Benefits



- Regularly monitoring for alignment issues fosters confidence in the integrity of the data, which is crucial for decision-making processes within the organization.
- Timely alerts allow for immediate investigation and resolution of data discrepancies, reducing the risk of prolonged issues and potential operational impacts.
- Notifications facilitate a more efficient reconciliation process

INTERVAL START	INTERVAL END	APPLICATION DATA OBJECT		
2025-05-16 07:00	2025-05-16 07:14	M3 MPLIND	Checksum verification	Instance count
			Correct	5 (Application) 5 (Data Lake)
			Row count	319239 (Application) [Play] (Compass)
2025-05-16 07:00	2025-05-16 07:14	M3 MITBAL		
2025-05-16 07:00	2025-05-16 07:14	M3 MPLINE		
2025-05-16 07:00	2025-05-16 07:14	M3 MPHEAD		

Support for rerunning reconciliation on multiple mismatches

Problem

- Admins couldn't easily rerun reconciliation on mismatch cards to check if late-arriving data resolved the issue. Manually clicking each card was inefficient at scale. The manual process compares the Infor Application's submitted reconciliation manifest with the current state of data in the Data Lake.

Solution

- A new 'Reconcile all mismatches' icon allows admins to rerun reconciliation for all mismatch cards at once, based on the selected ingestion method.
- Once triggered, affected cards will show a 'Running' status during the process.

Benefit

- This streamlines manual bulk reconciliation, improves user experience, and helps quickly detect resolved mismatches

100288

Data Fabric Data Ledger

KB2105703

of mismatched cards is currently in progress. Refresh the list of cards to check for updates.

INTERVAL START	INTERVAL END	APPLICATION
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2
2/3/25, 9:15 AM	2/3/25, 9:20 AM	LOADTEST_OLD_2

Health Service notifications on Data Ledger reconciliation mismatch

Problem

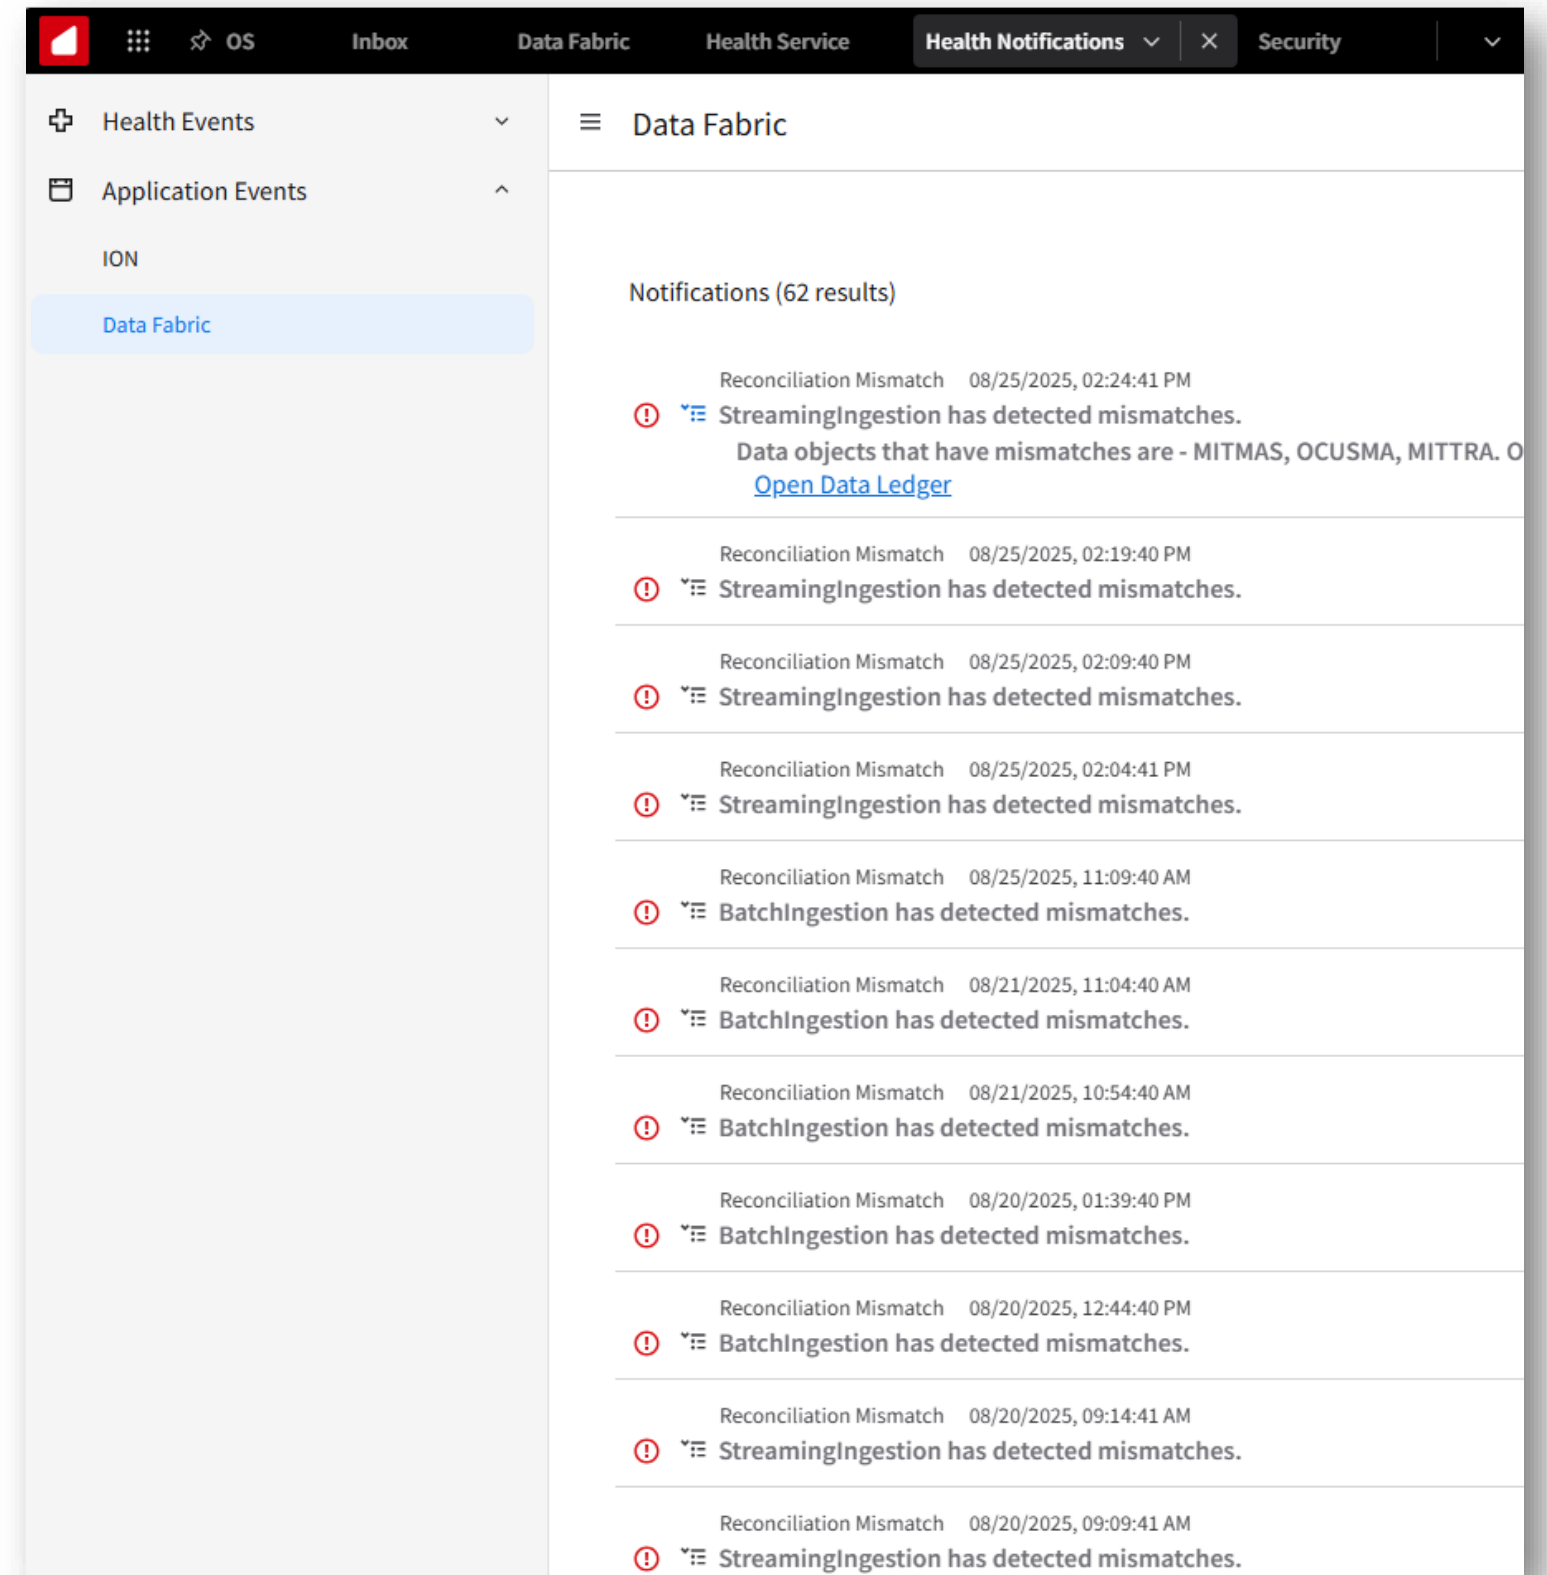
- Admins had to manually check the Data Ledger page to detect discrepancies between the system of record and the Data Lake. They preferred automated notifications

Solution

- Data Ledger now uses Health Service to send 'Reconciliation Mismatch' events when automatic reconciliation detects a mismatch.
- Notifications are triggered approximately 15 minutes after the Infor Application submits the reconciliation manifest to the Data Ledger service.
- To use Health Service with Data Ledger users need the **DATAFABRIC-HealthServiceAdmin** IFS Security role.
- Events are sent for BatchIngestion or StreamingIngestion, depending on the Infor Application and ingestion method.
- Refer to the Health Service user guide to set up email or Portal notifications.

Benefit

- Regular monitoring and timely alerts help ensure data integrity and enable quick resolution of discrepancies, while notifications keep admins informed and streamline reconciliation.



Agenda

1

M3 Publishing data

Data Lake Publisher enhancements

- Table groups and content management
- Initial load or partial reload managed and controlled within DLP.

2

Data Reconciliation

Data Ledger principles for validating consistency in batched and streamed data in Data Lake

- Streamed data
- ION ingestion
- Batch ingested data

3

Data Fabric

Last year's enhancements and news in 2025.10

- Data Catalog
- Compass
- Stream Pipelines
- Undo Purge

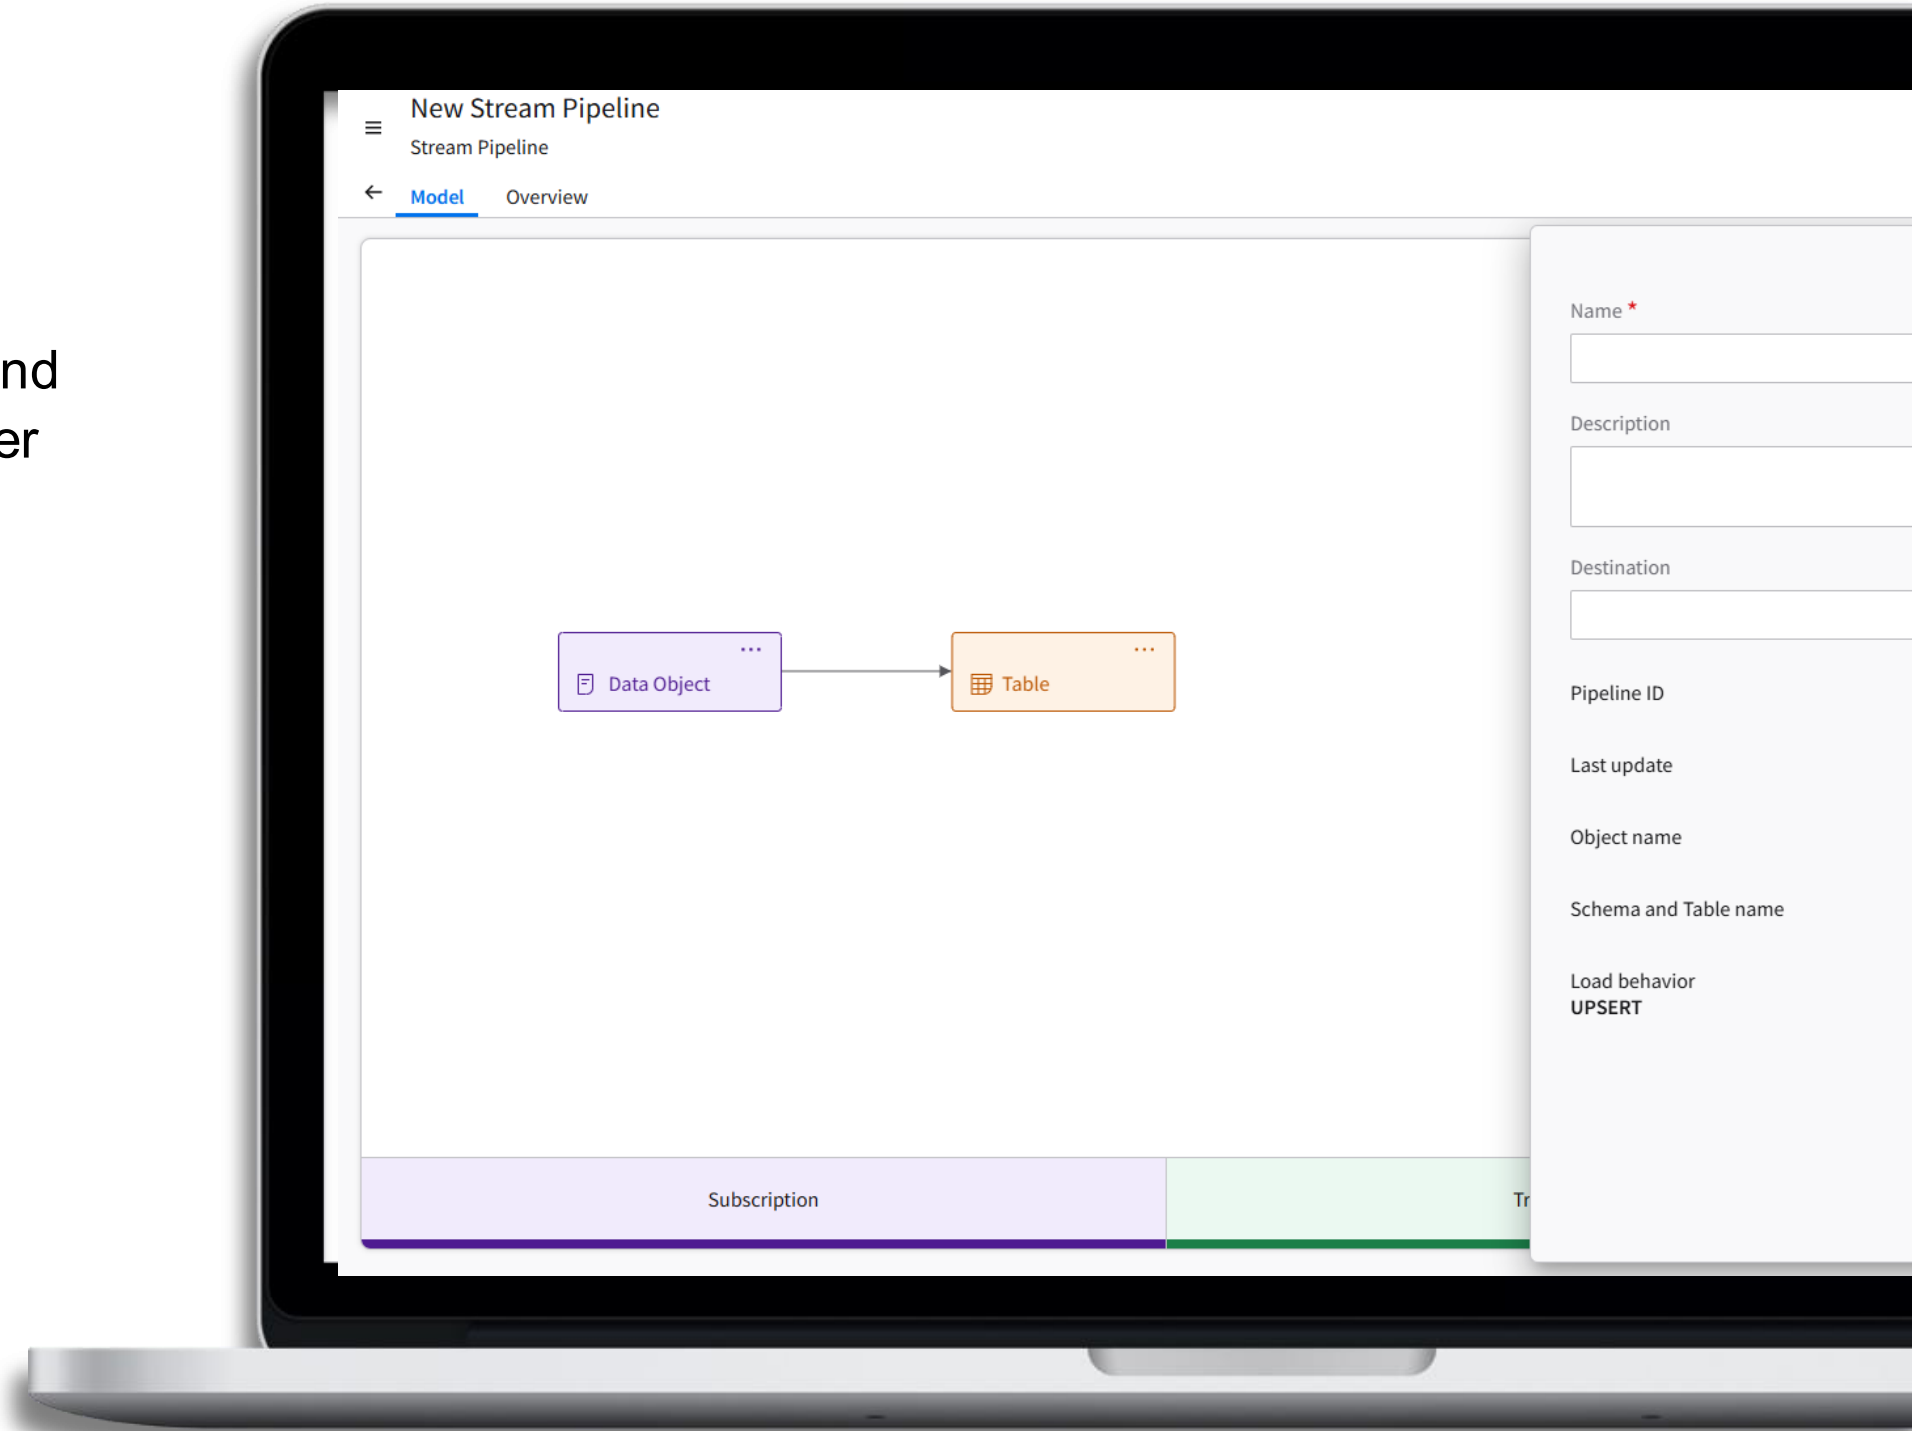
4

Source Data Access and Q&A

- Future features in development.
- Questions and Answers

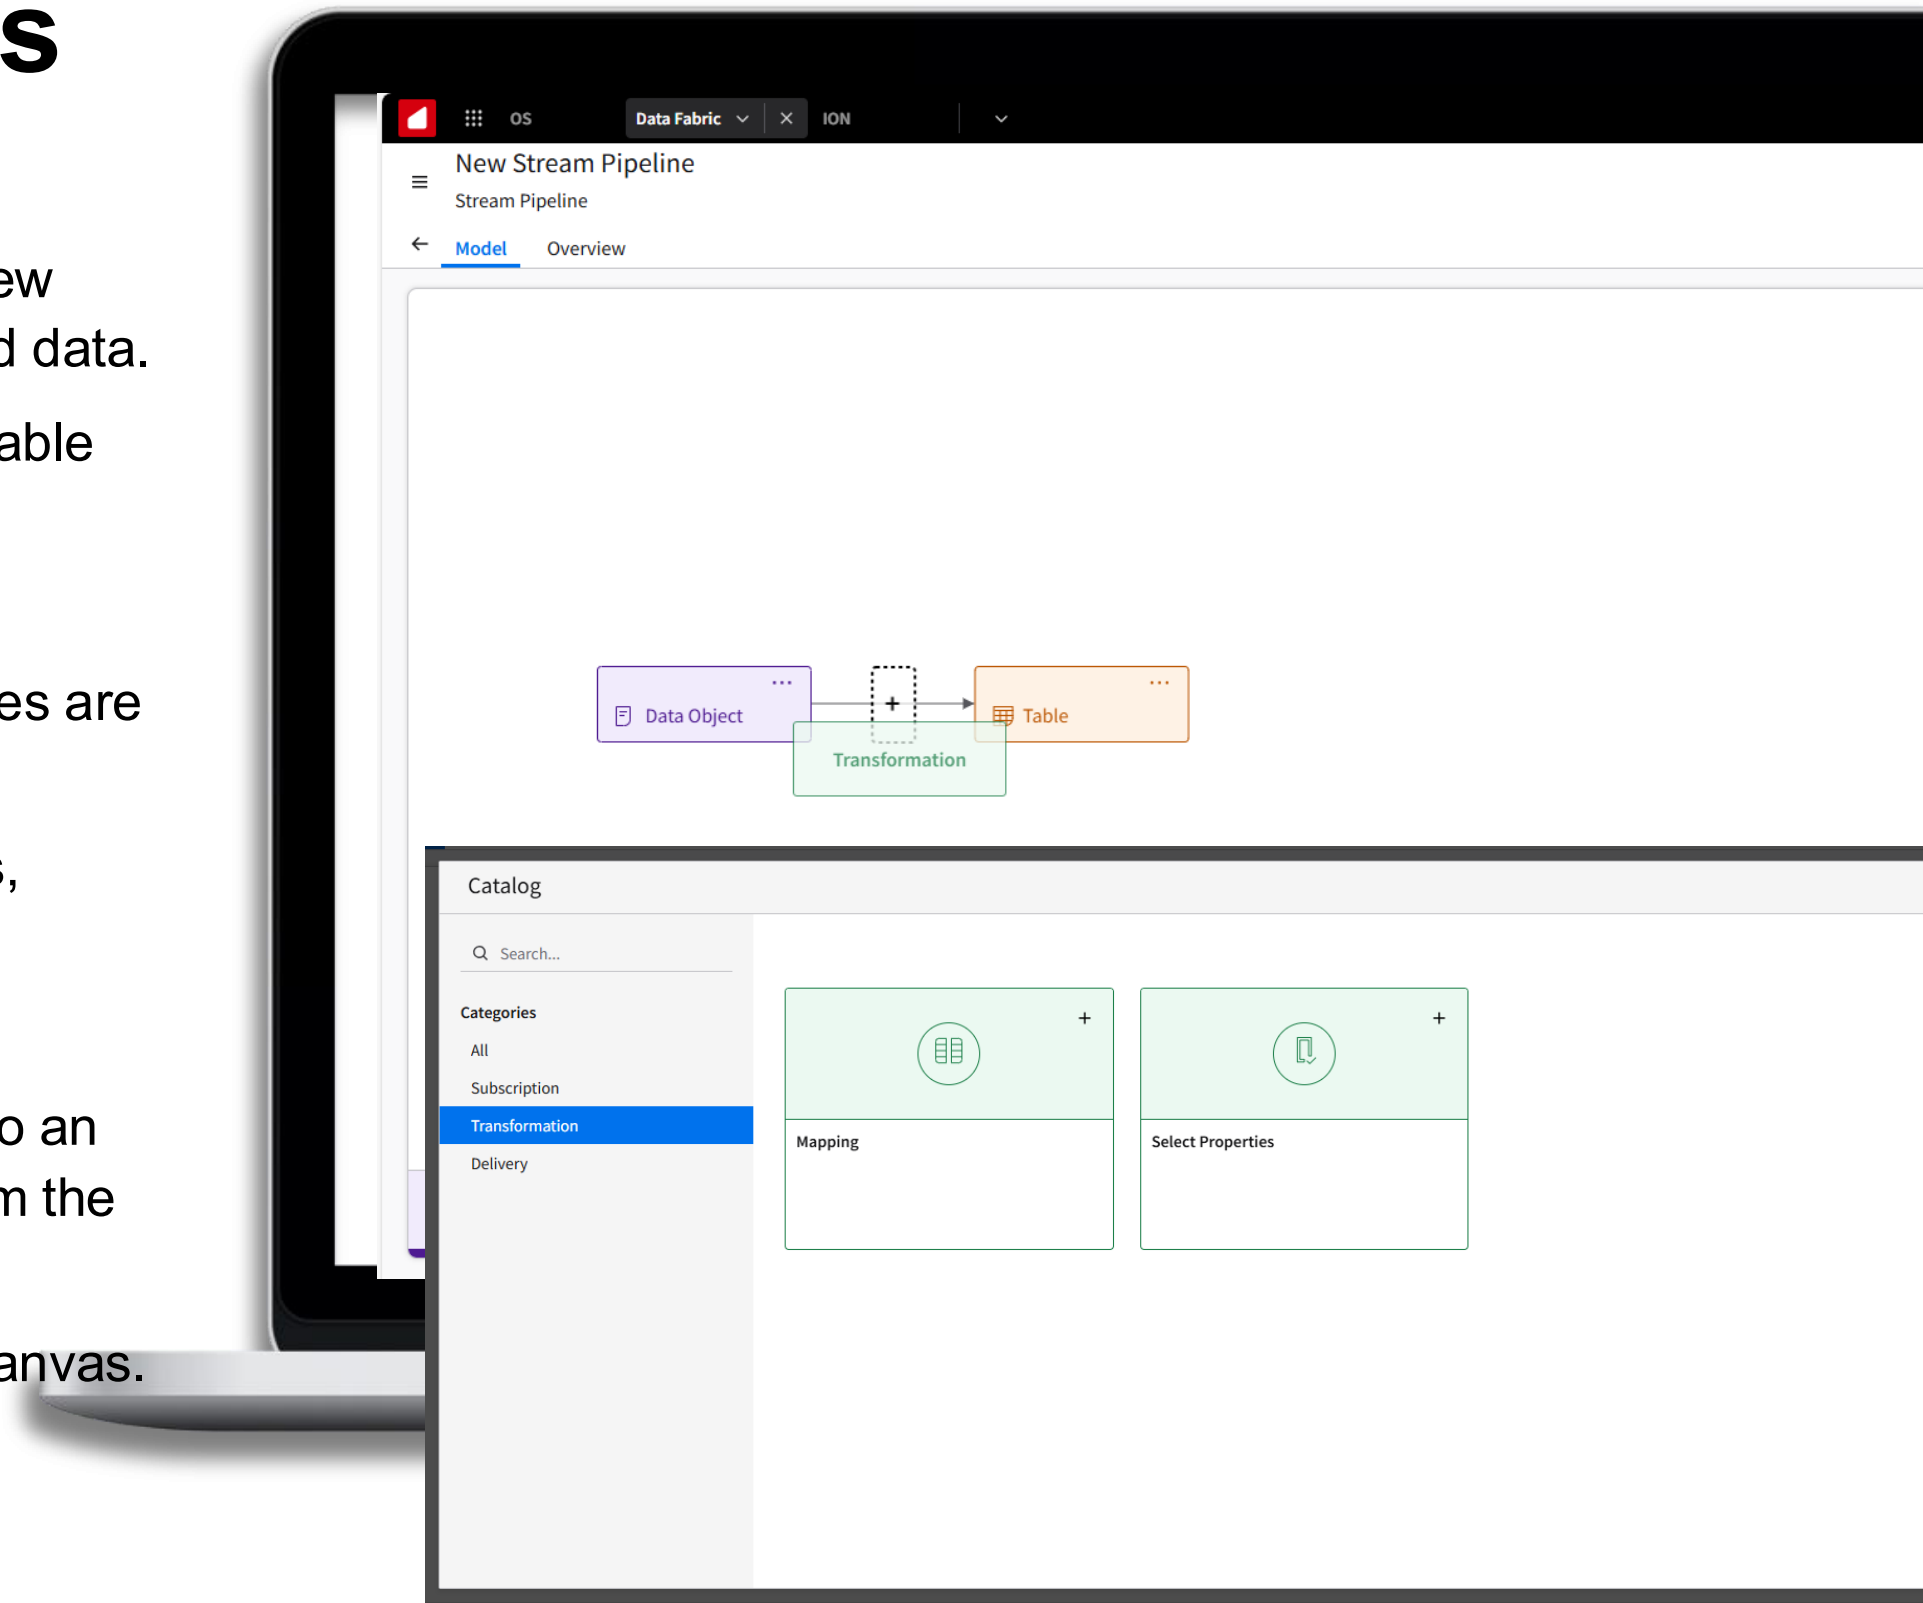
Creating a new pipeline with default activities

- When a new pipeline is created, the required Object and Table activities are automatically placed on the modeler canvas.



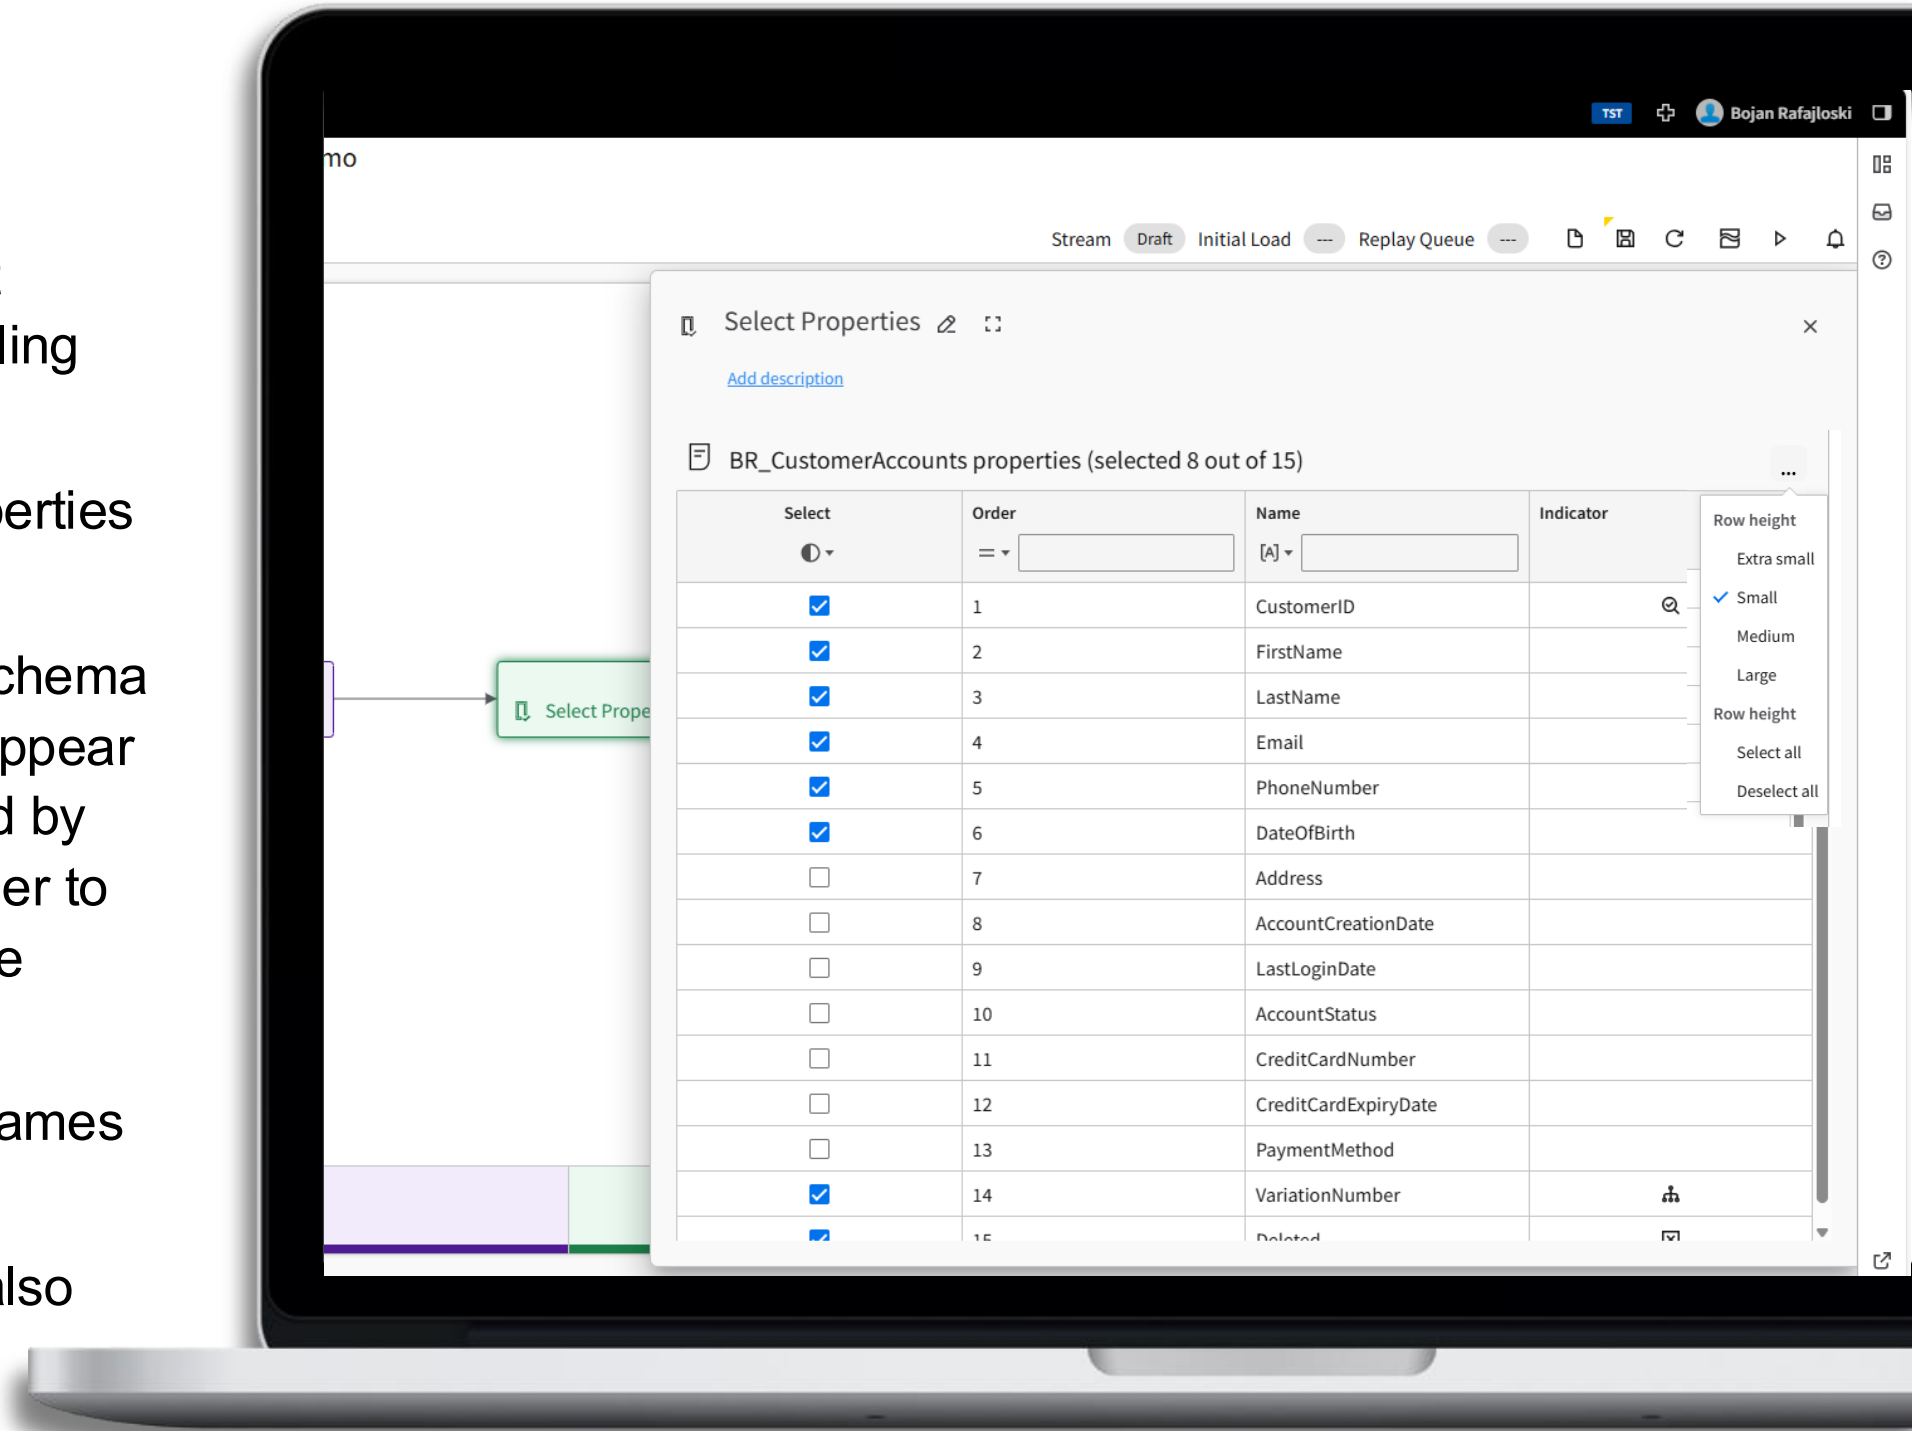
Transformation activities

- The Stream Pipelines activity catalog now includes new Transformation activities designed to modify streamed data.
- These activities are placed between the Object and Table activities.
- The available transformations are:
 - **Select Properties** – Choose which object properties are streamed to the destination table.
 - **Mapping** – Map object properties to table columns, supporting custom naming and partial matches.
- To add Transformation activity
 - Option 1: Drag & drop the Transformation activity to an empty place in the canvas. Choose the activity from the catalog menu.
 - Option 2: Hover & click on an empty place in the canvas. Choose the activity from the catalog menu.



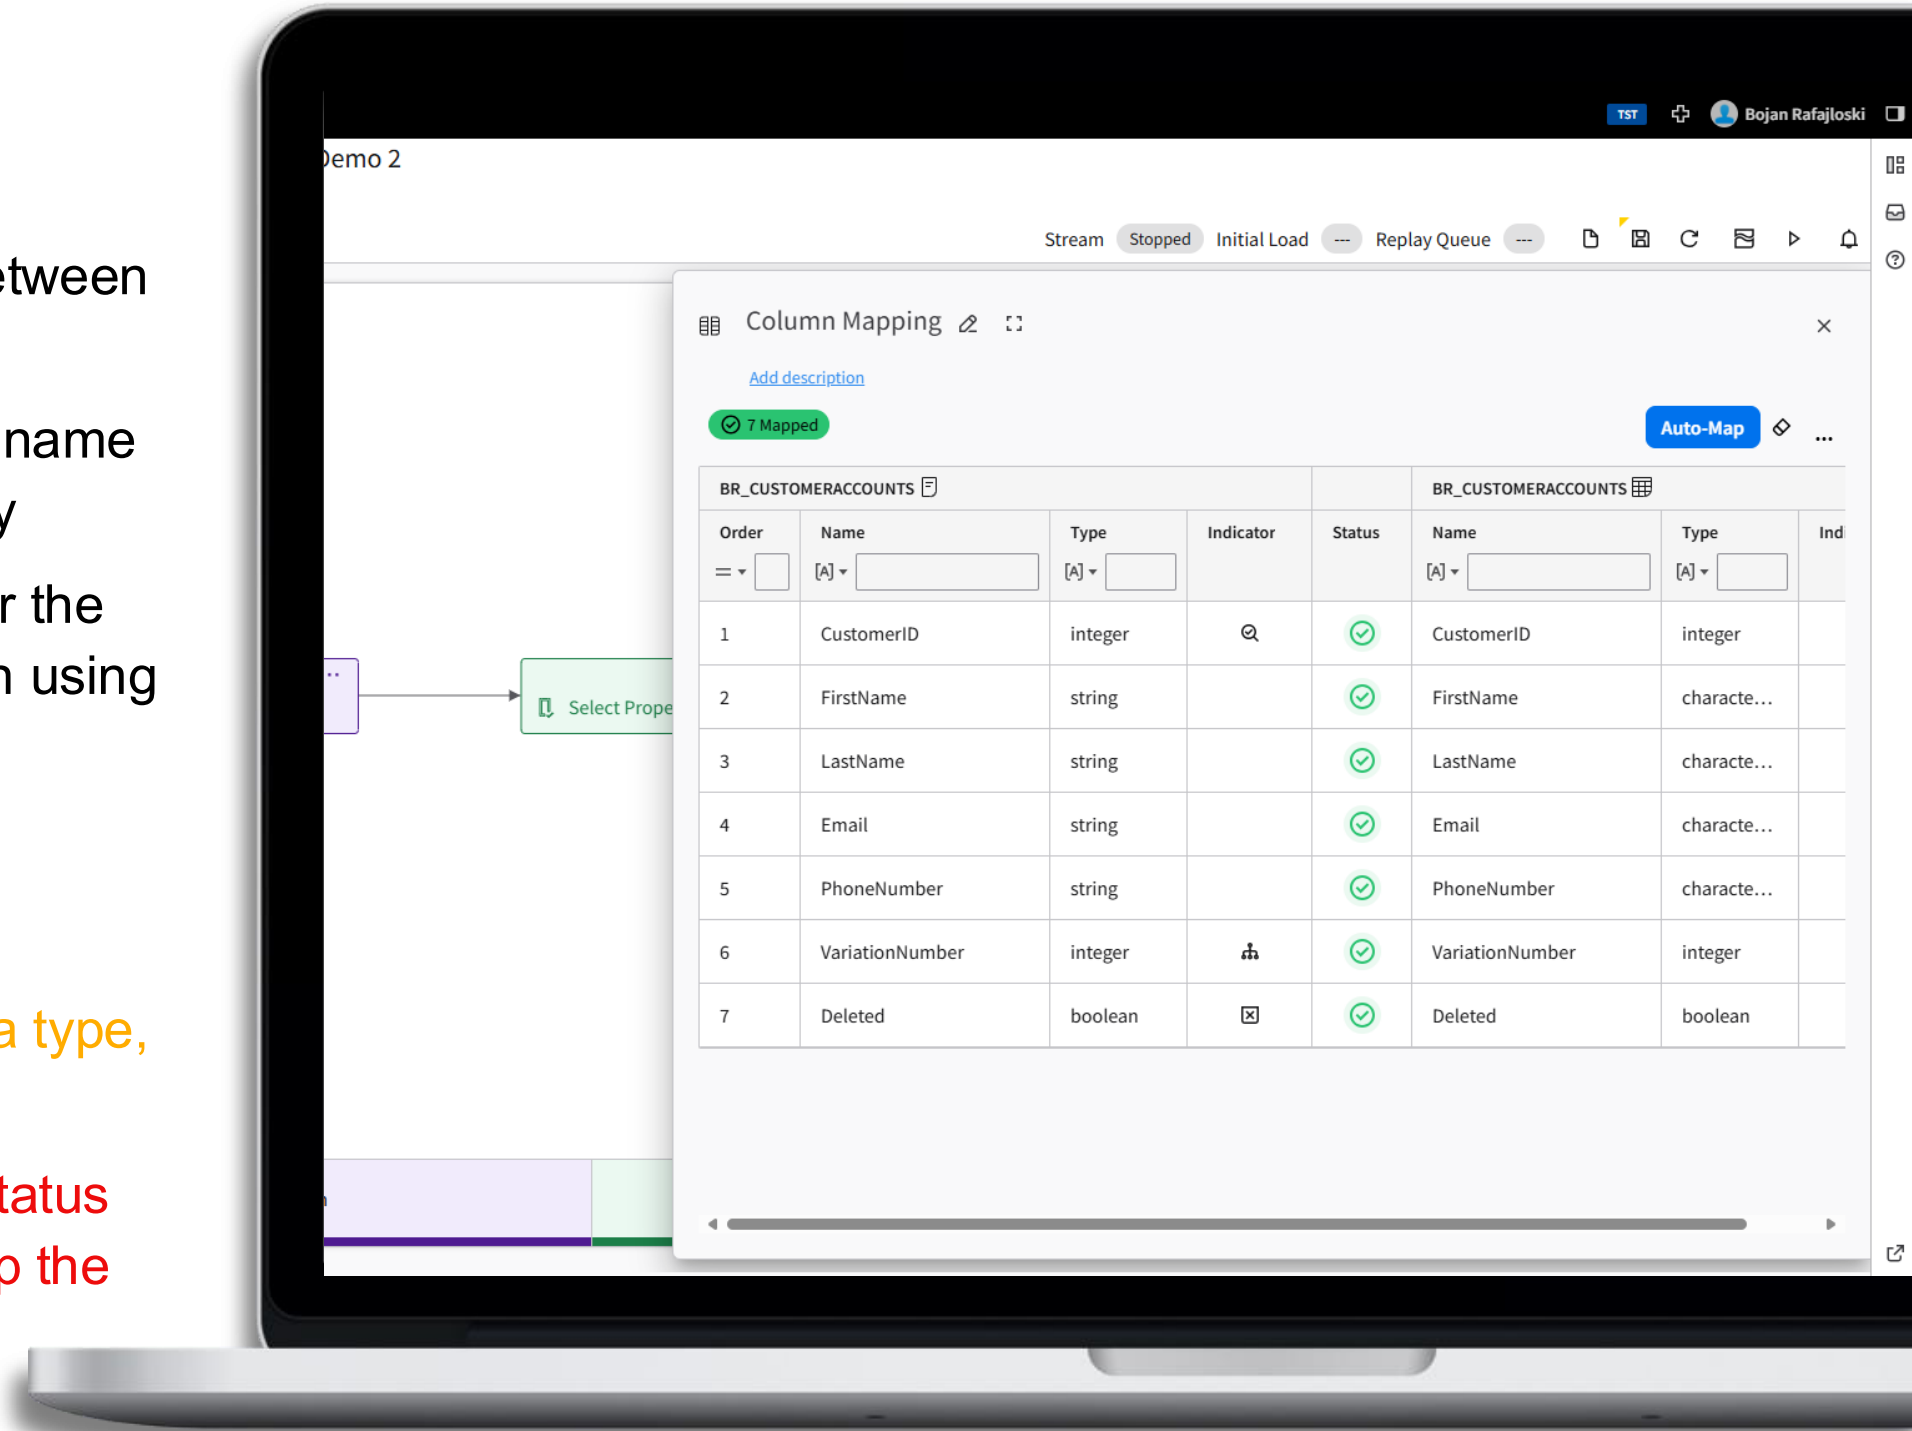
Select properties

- Select Properties allows users to choose which object properties are streamed to the destination table, enabling them to exclude columns that are not needed.
- Using checkboxes, users can easily select which properties should proceed to the next step in the pipeline.
- If the pipeline is activated or stopped and the object schema is updated with new properties, those properties will appear in the Select Properties activity but will not be selected by default. This gives users control over when and whether to include new properties in the stream, allowing for more deliberate and flexible schema management.
- Using the column filters you can search for property names and types
- In the '...' menu you can select and deselect all, and also change the row-height for the UI datagrid



Mapping

- Mapping activity allows you to choose the mapping between the Object properties and columns from the Table.
- Using Auto-map button, it will try to detect the column name that match the property name either exactly or partially
- To manually choose a column, in the Name column for the Table, click on an empty name and choose the column using the select dropdown
- The validation will show the following status:
 - Properties that are correctly mapped to a column
 - Properties mapped to a column with a different data type, from what is defined in the object schema
 - Properties that are not mapped to a column. This status will prevent pipeline activation, requiring you to map the property to a column, or alternatively to remove the property with Select Properties activity
- You can clear the select columns using the Eraser icon.



Best Practice

- Even though they are optional, it is recommended to use Select Properties and Mapping activities in your pipelines.
- It allows for efficient control over object schema changes and offer flexibility in column naming - so you're no longer restricted to including every property and using identical names or letter casing.

Restoring purged data objects

Problem

- Data Lake objects are permanently purged, and there is no option to reverse this action if it occurs accidentally or due to malicious intent.

Solution

- Go to the Data Fabric Purge page.
- Select a purge event log to restore its deleted data objects.
- The restoration process may take time depending on the number of objects.
- While restoring, the purge event show a status of 'Restoring', once completed, the status will update to 'Restored'. Restoring process cannot be stopped or reverted.
- Restored objects are logged under the same purge event, including who restored and when.
- Compass cached data will be cleared to ensure restored objects are re-cached in future queries.
- Purge events older than 90 days cannot be restored.

Benefit

- Safety net against accidental deletions or malicious activities
- Minimizing potential disruptions to business operations and ensuring that valuable data assets are not permanently lost

The screenshot displays the Data Fabric Purge interface. At the top, there are filters for 'Date range To' (with a date picker showing 'M/d/yyyy h:mm:ss a') and 'Data objects' (a dropdown menu). Below these is a 'Location' filter. The main area contains a table with columns for 'Restored Start Date', 'Restored End Date', and 'Restored by'. The table has a header row with date pickers and a user selection dropdown, and one data row showing a restoration event on 8/25/25. The footer of the table shows 'of 131' records and '25 records per page'.

	Restored Start Date	Restored End Date	Restored by
	M/d/yyyy h:mm:ss a	M/d/yyyy h:mm:ss a	[A]
	8/25/25, 12:25:19 PM	8/25/25, 12:25:33 PM	Sindhura Kavali

Data Catalog application migrated from ION to Data Fabric

Problem

- Data Fabric admins wanted to access Data Catalog directly from the Data Fabric application, without needing access to ION Desk.

Solution

- The Data Catalog application and APIs have been moved from ION to Data Fabric.
- The application will remain available in ION for 1 year, after which it will only exist in Data Fabric.
- ION based Data Catalog APIs will be deprecated and existing integrations should switch to the Data Fabric APIs.
- IFS Role permissions have been migrated, so existing ION Data Catalog users will automatically have access to Data Catalog in Data Fabric.

Benefit

- This tighter integration enhances metadata and data platform capabilities, offering better tools for exploration and analysis.

100288

Data Fabric Data Ledger

KB2105703

The screenshot shows the Data Fabric application interface. The top navigation bar includes 'Data Fabric', 'Health Service', 'Health Notifications', and 'Security'. The left sidebar contains a navigation menu with items like 'Home', 'Data Lake', 'Data Warehouse', 'Compass', 'Data Catalog', 'Object Schemas', 'Data Lake Views', 'Data Warehouse Schemas', 'Schema Extensions', 'Locale Selections', 'Metagraphs', 'Graph', 'Pipelines', 'Data Orchestrator', 'Security', 'Utilities', 'Configuration Management', and 'About Data Fabric'. The main content area is titled 'Object Schemas' and features a search bar with the text 'LN_'. Below the search bar, there are filters for 'Library', 'Object Type', 'Last updated on', and 'Last updated by'. A 'Search' button is located at the bottom of the filter section. The main content area displays a list of object schemas with the following details:

Name	Type	Subtype	Last updated	Last updated by
LN_tdlsa001	JSON	Newline-delimited	2023 Oct 3	Rajwant Kaur
LN_tfgld172_2	DSV	CSV	2023 Jul 25	Tomas Vlk
LN_tfgld170_textt	DSV	CSV	2023 Jul 24	Tomas Vlk
LN_ticst002	JSON	Newline-delimited	2023 Jul 21	Mike Kalinowski
LN_qmptc101	JSON	Newline-delimited	2022 May ?	

Compass /results/ API support for Accept-Encoding: gzip

Problem

- Not supporting gzip compression can make integration harder for systems that prefer it over Deflate. Various tools and APIs use gzip by default, so lack of support may lead to extra configuration, reduced compatibility, and slower performance.

Solution

- **GET /DATAFABRIC/compass/v2/jobs/{queryId}/result/** now Accept-Encoding supports returning content in **gzip** compression

Benefit

- Supporting gzip compression improves compatibility with tools and APIs that use it by default or have challenges using deflate, reducing need for custom solutions

Health Notifications Security API Gateway

Submits a Compass SQL job request.

{queryId}/cancel/ Cancels running SQL query.

{queryId}/result/ Get the result for a queryId that was provided from a Compass job.

Results using the queryId that was provided when submitting the a query.
Compass Status API indicates that the query has finished successfully, failed, or has been cancelled. A query that finished successfully returns the message.

Description

The unique queryId provided as a response by the Compass Jobs API when the query was submitted.

queryId

Offset is the row offset for the result; offset is 0-based.

0

Limit is the maximum number of rows in the result. The limit is 100 000 rows max with a maximum 10MB size cap.

0

Specifies the returned content-encoding. The supported encodings are gzip, deflate, and identity.

--

gzip

deflate

identity

Execute

Description

The method returns in JSON, NDJSON, or CSV format.

ETL Tool installation file is available in Downloads page

Problem

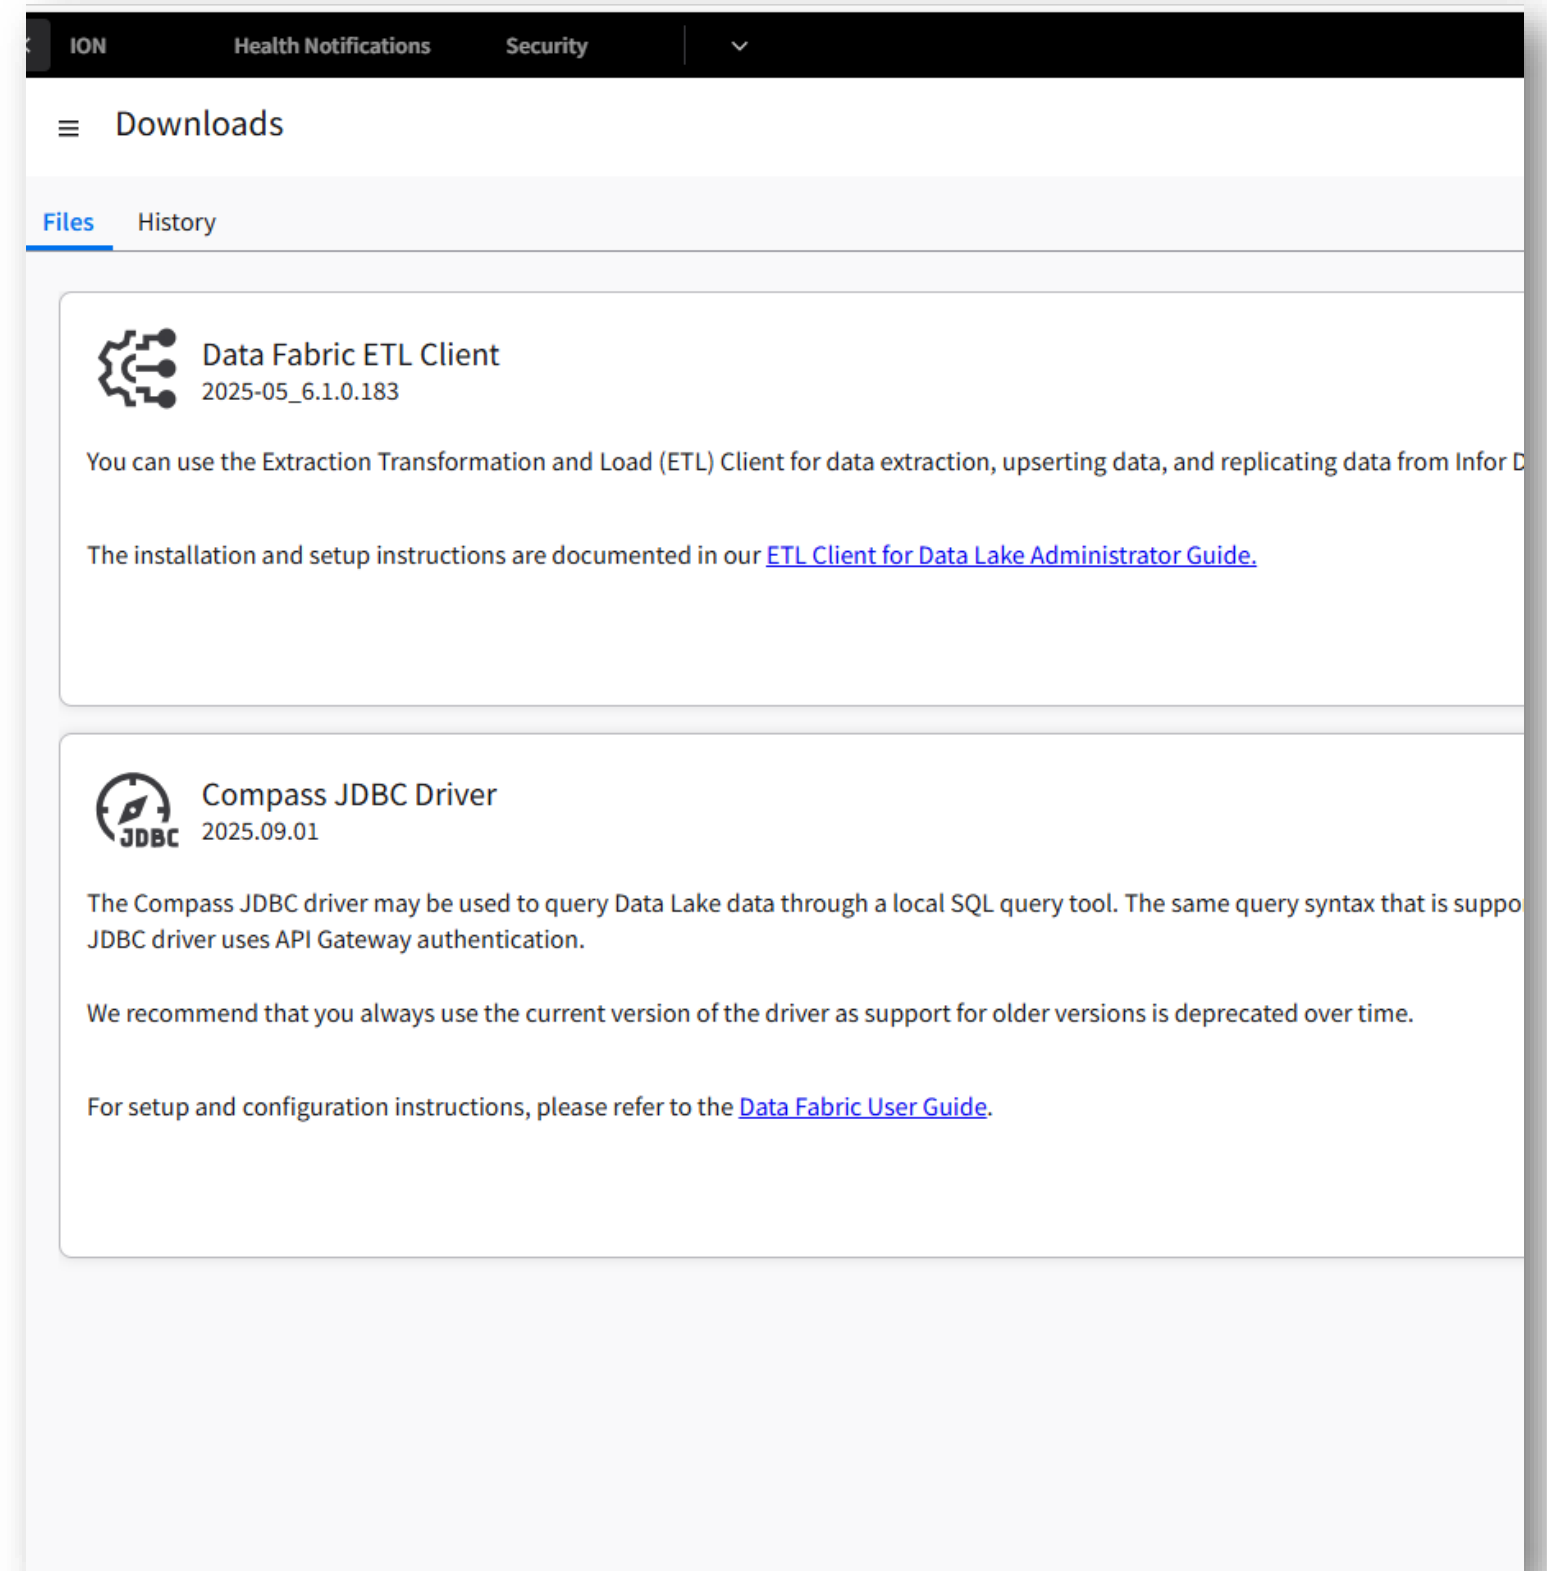
- Previously, the ETL Tool was only available through ION, making access less convenient for Data Fabric users.

Solution

- The ETL Tool is now available directly in Data Fabric under Utilities → Downloads, eliminating the need to access ION.

Benefit

- This improves accessibility and streamlines workflows by allowing admins to download and use the ETL Tool directly within the Data Fabric environment.



Data Catalog enhancements

2024

- Check for duplicate metadata when registered via API
- Configuration Management Local Packages - Improvement for schema extensions

2025.04

- Support for automatic schema updates in Compass

2025.10

- Migrate Data Catalog APIs to the Data Fabric API Gateway suite
- Move to Data Fabric
- Configuration Management - Global package support

Security and ConfigMgmt

2024

- Data Fabric Security now supports the ability to manage access to data stored in Data Lake.
- For new tenants, security roles provisioned for Data Fabric will be listed on the Data Fabric Security page.
- **Configuration Management Tags** A new page, called Configuration Management > Tags, has been added to Data Fabric. This page presents a comprehensive list of tags that can be utilized to link your Data Fabric resources, aiding you in the preparation of a local package.
- Configuration Management Tags - Resource security
Access controls are now used to determine what a user can see when viewing the details of a tag in Data Fabric.

2025.04

- Navigation menu reorganization
- Data Fabric Security: Change to default sort order on the Data Security tab
- Data Fabric Security: Support the ability to import and export security profiles
- Data Fabric: Downloads page

Data Ledger

2024

- Data Lake: Atlas data grid filter enhancements

2025

- Data Ledger: Enhancement on event reconciliation for streamed data
- Data Ledger: Rerun the reconciliation for all mismatch cards

2025.10

- Data Ledger: Infor OS Health Services adoption
- Data Lake - Purge Undo

Stream Pipelines

2024

- Snowflake Destination
- Configuration Management: Local package support
- Destinations - Rename the PostgreSQL destinations for consistency
- Remove the requirement to have a default warehouse for Snowflake user
- Support mapping Number to other data types in Snowflake
- SQL Server Destination

2025.06

- Stream pipelines - Optimize connection retry handling
- Stream pipelines - UI enhancement - Display pipelineId and other details

2025.10

- Stream pipelines - Default activities when creating a pipeline
- Stream pipelines - Mapping source properties to target table columns with different names
- Stream pipelines - Select object properties (columns) to stream

April 2025

Stream Pipelines and Data Lake

- Stream Pipeline destinations: RDS SQL Server / Azure SQL
- Batch time to Data Lake when streaming reduced from 10 to 5 minutes

Compass

- Compass query result 100 < 10 000 rows
- Compass SQL performance optimization for Analytical mode
- Synchronized schema updates across Data Catalog & Compass SQL

Data Ledger

- Data Ledger – refresh multiple mismatches
- Data Ledger enhancements for streaming

Data Security Import/Export

Summary

- How to restore purged data objects from the Purge module.
- How Health Service sends notifications when Data Ledger reconciliation mismatches occur.
- How to rerun reconciliation on multiple mismatches in Data Ledger using the 'Reconcile all mismatches' feature (released in v2025.07).
- That the Data Catalog application has been migrated from ION to Data Fabric, with updated access and API usage.
- That the Compass API now supports gzip compression via the Accept-Encoding: gzip header (released in v2025.06).
- That the ETL Tool installation file is now available on the Downloads page.

Agenda

1

M3 Publishing data

Data Lake Publisher enhancements

- Table groups and content management
- Initial load or partial reload managed and controlled within DLP.

2

Data Reconciliation

Data Ledger principles for validating consistency in batched and streamed data in Data Lake

- Streamed data
- ION ingestion
- Batch ingested data

3

Data Fabric

Last year's enhancements and news in 2025.10

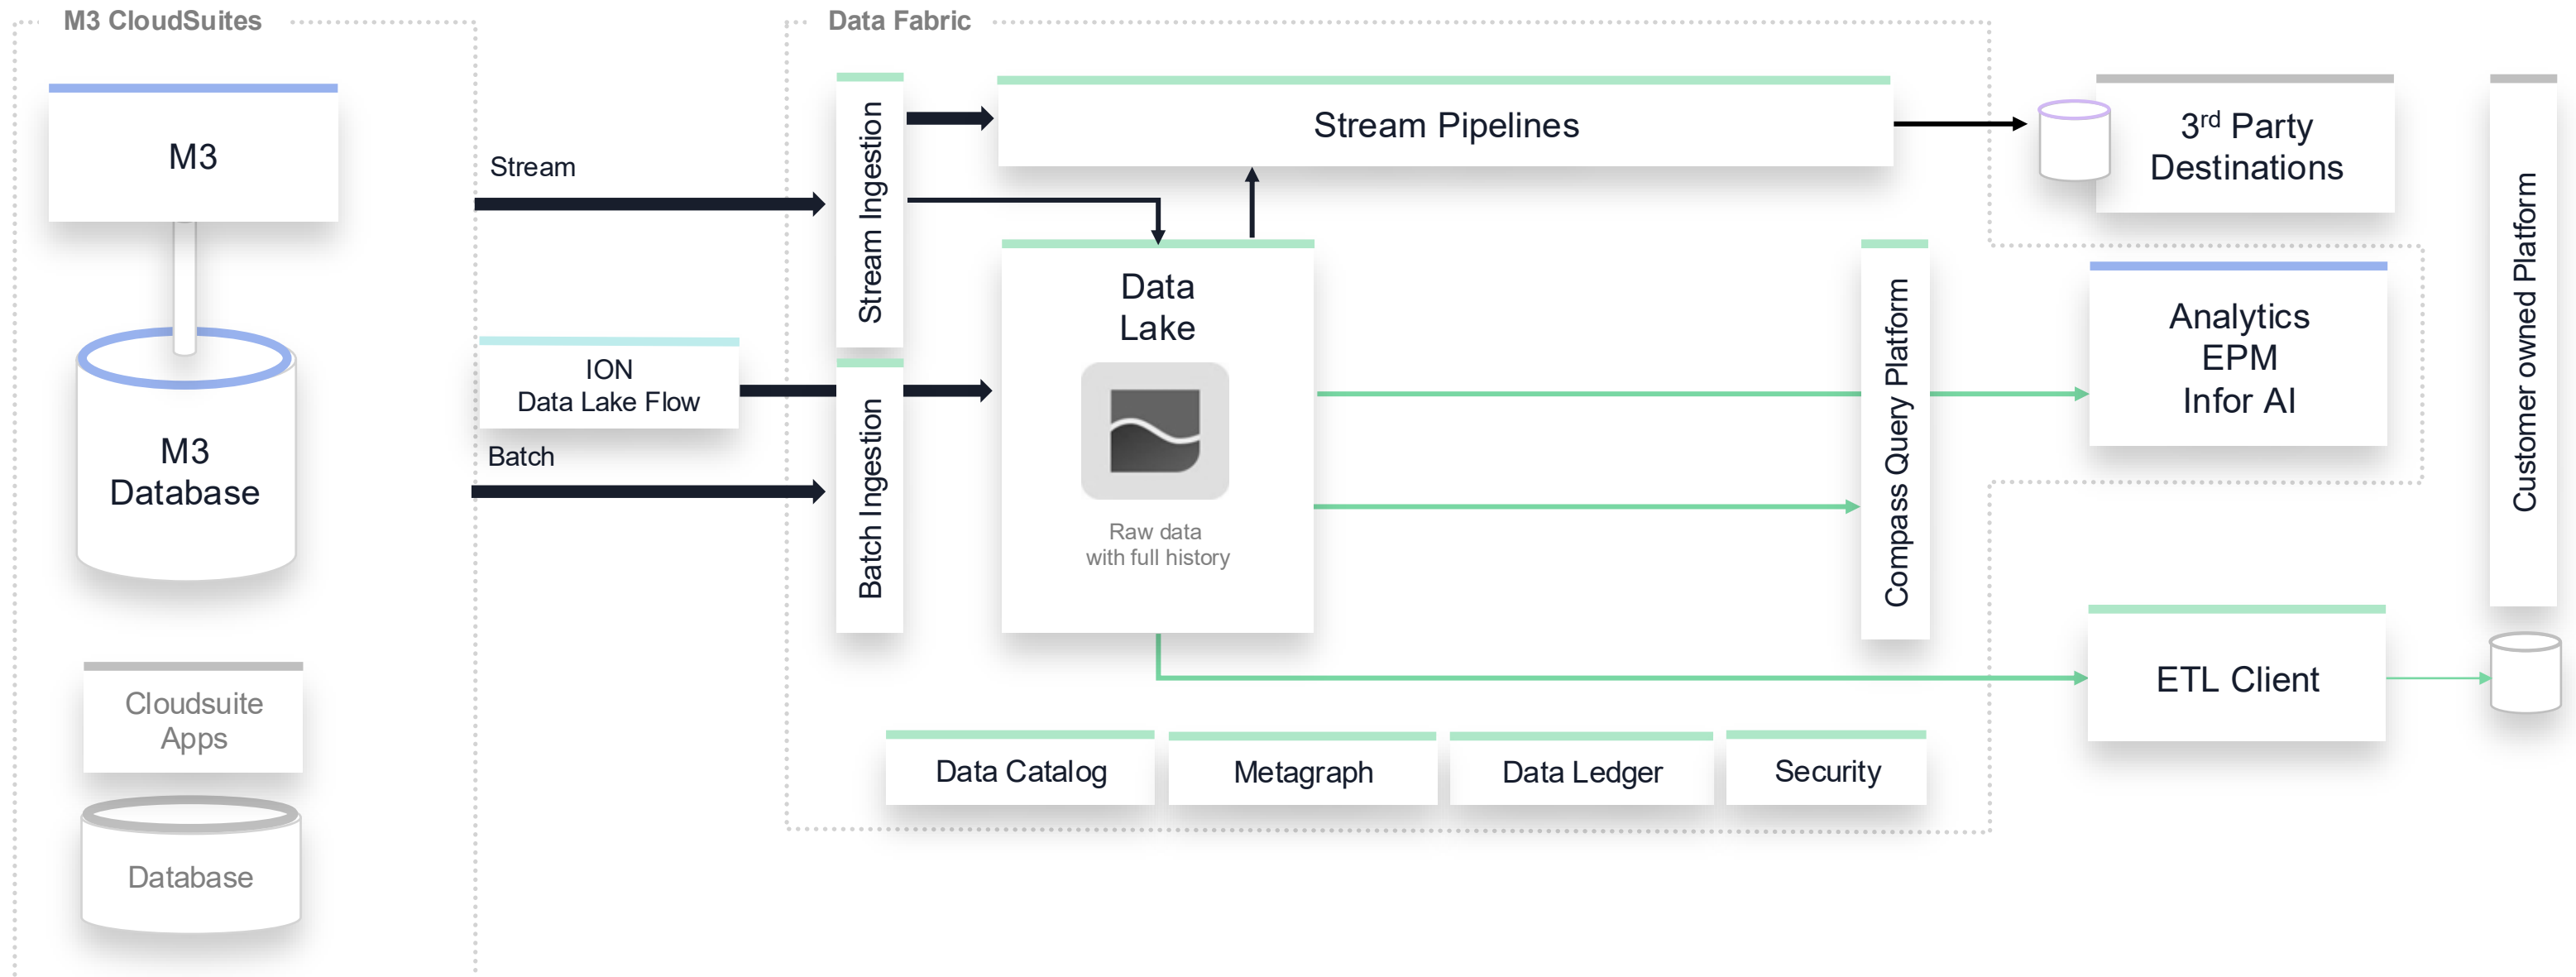
- Data Catalog
- Compass
- Stream Pipelines
- Undo Purge

4

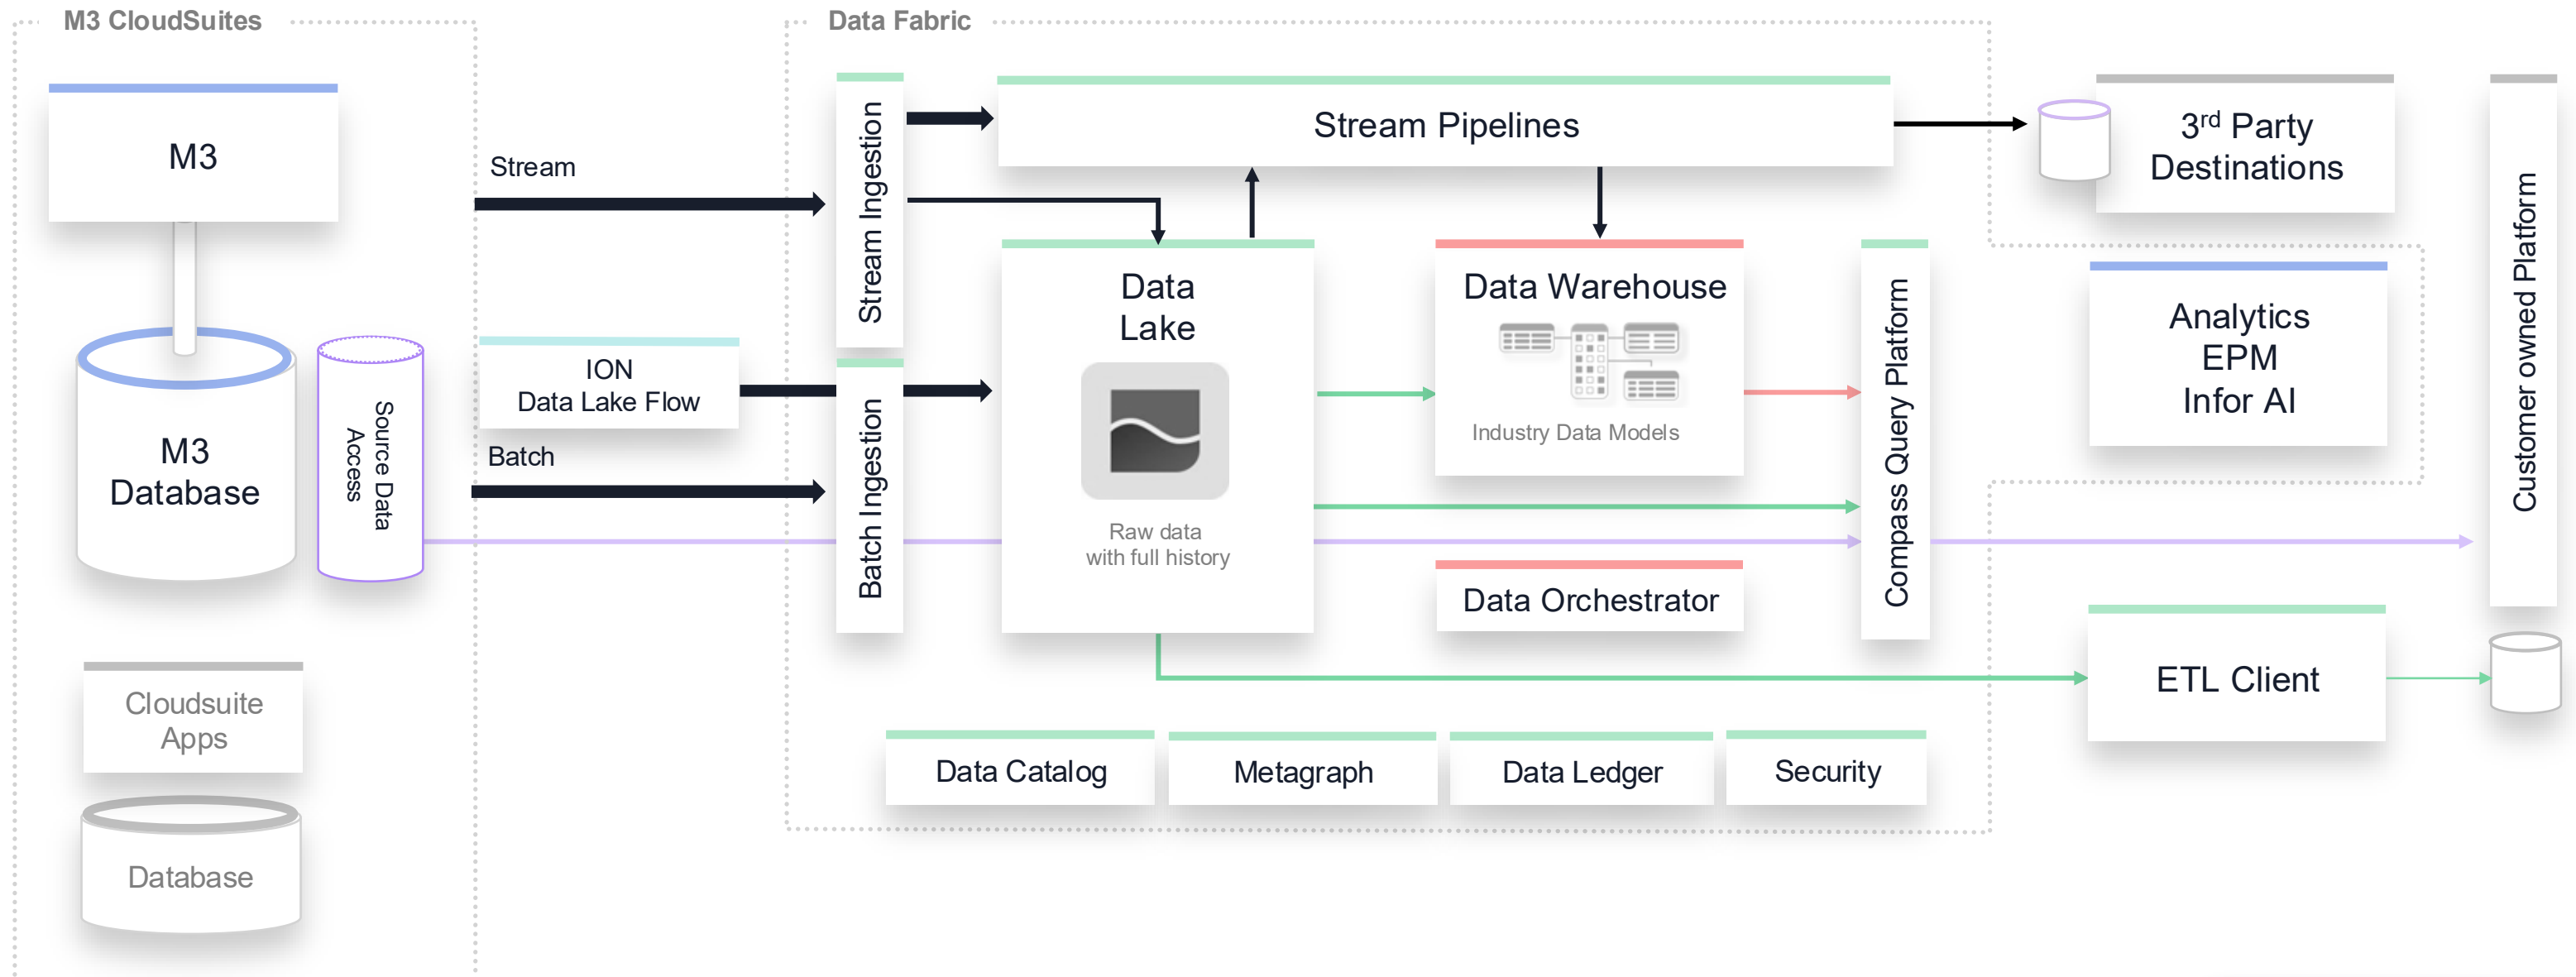
Source Data Access and Q&A

- Future features in development.
- Questions and Answers

M3 and Data Fabric 2025.10

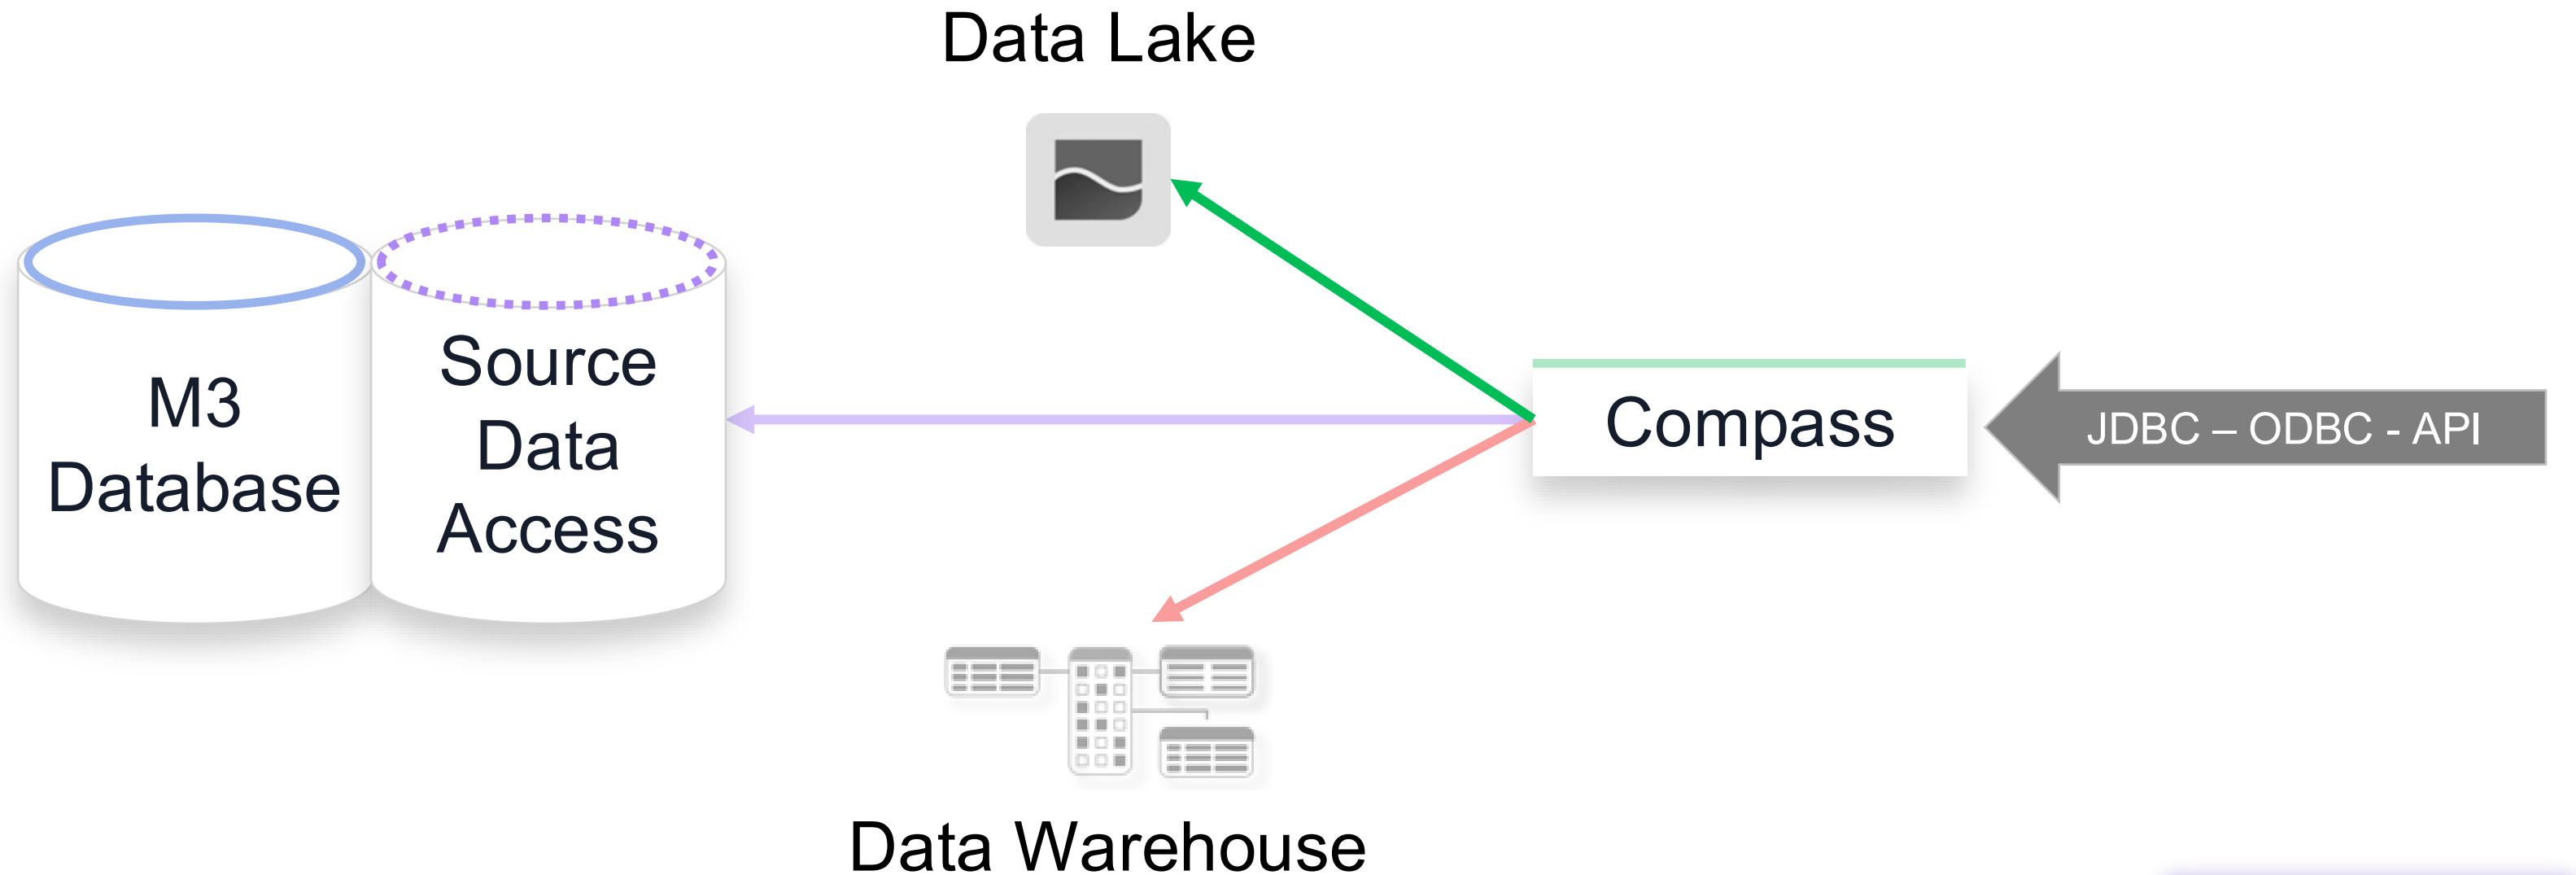


M3 and Source Data Access



Source Data Access
Limited Available
October 2025

M3 and Source Data Access



Source Data Access
Limited Available
October 2025

Data Repositories and Data Access



Data Lake

A highly scalable big data storage layer for replicated ERP datasets, enabling historical traceability and complex analytics.

Ideal for:

- Big data storage with full historical transaction store
- Batch scheduled data extractions and analytical queries
- Cost-effective long-term data retention

Limitations:

- Not optimized for real-time or performance critical integrations
- Not ideal for high-volume query execution



Data Warehouse

Data Warehouse is a relational database acting as a central hub for curated data. It allows to load, transform, and build curated datasets for analytical and reporting purposes.

Ideal for:

- Curated and prebuilt data models
- High volume query execution
- Long running query workloads
- Complex reporting and BI dashboards
- Real-time streaming updates

Limitations:

- Higher cost for long-term data storage compared to Data Lake



Source Data Access

Provides read-only, direct access to the live CloudSuite ERP database, serving as the foundation for real-time transactional reporting and low-latency queries

Ideal for:

- Low-latency, performance-sensitive integrations
- Real-time transactional data
- Lightweight, ad hoc queries for timely decision-making

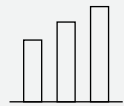
Limitations:

- Not suitable for complex joins or analytical workloads and high-volume query workloads
- Not suitable for continuous data extraction
- Requires deep understanding of ERP schema complexity

Executive Overview

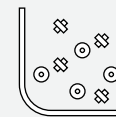
SDA delivers governed, read-only SQL access to live M3 data without replication or production impact

Purpose



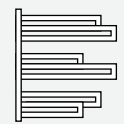
- Real-time visibility into live M3 data
- Targeted read-only access

Scope



- Validation, reconciliation, investigation
- Data quality, process checks
- Post-migration assurance, audit support

Foundation



- Delivered via Infor OS + Compass
- Authenticated, audited, governed
- Read-node guardrails

Value



- Faster root cause analysis
- Higher confidence in operational data
- Transparency without security/performance trade-offs

What Is Source Data Access



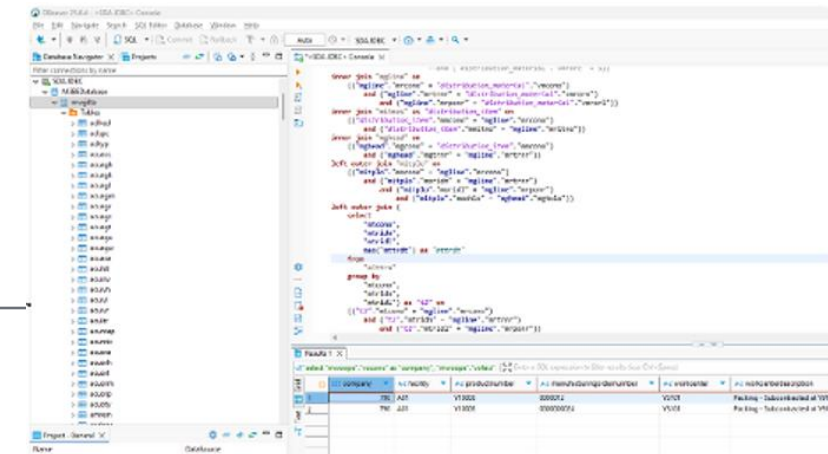
M3



Source Data Access (SDA)



Compass (secure gateway)



M3 Database
Live transactions
Protected system
No direct query access

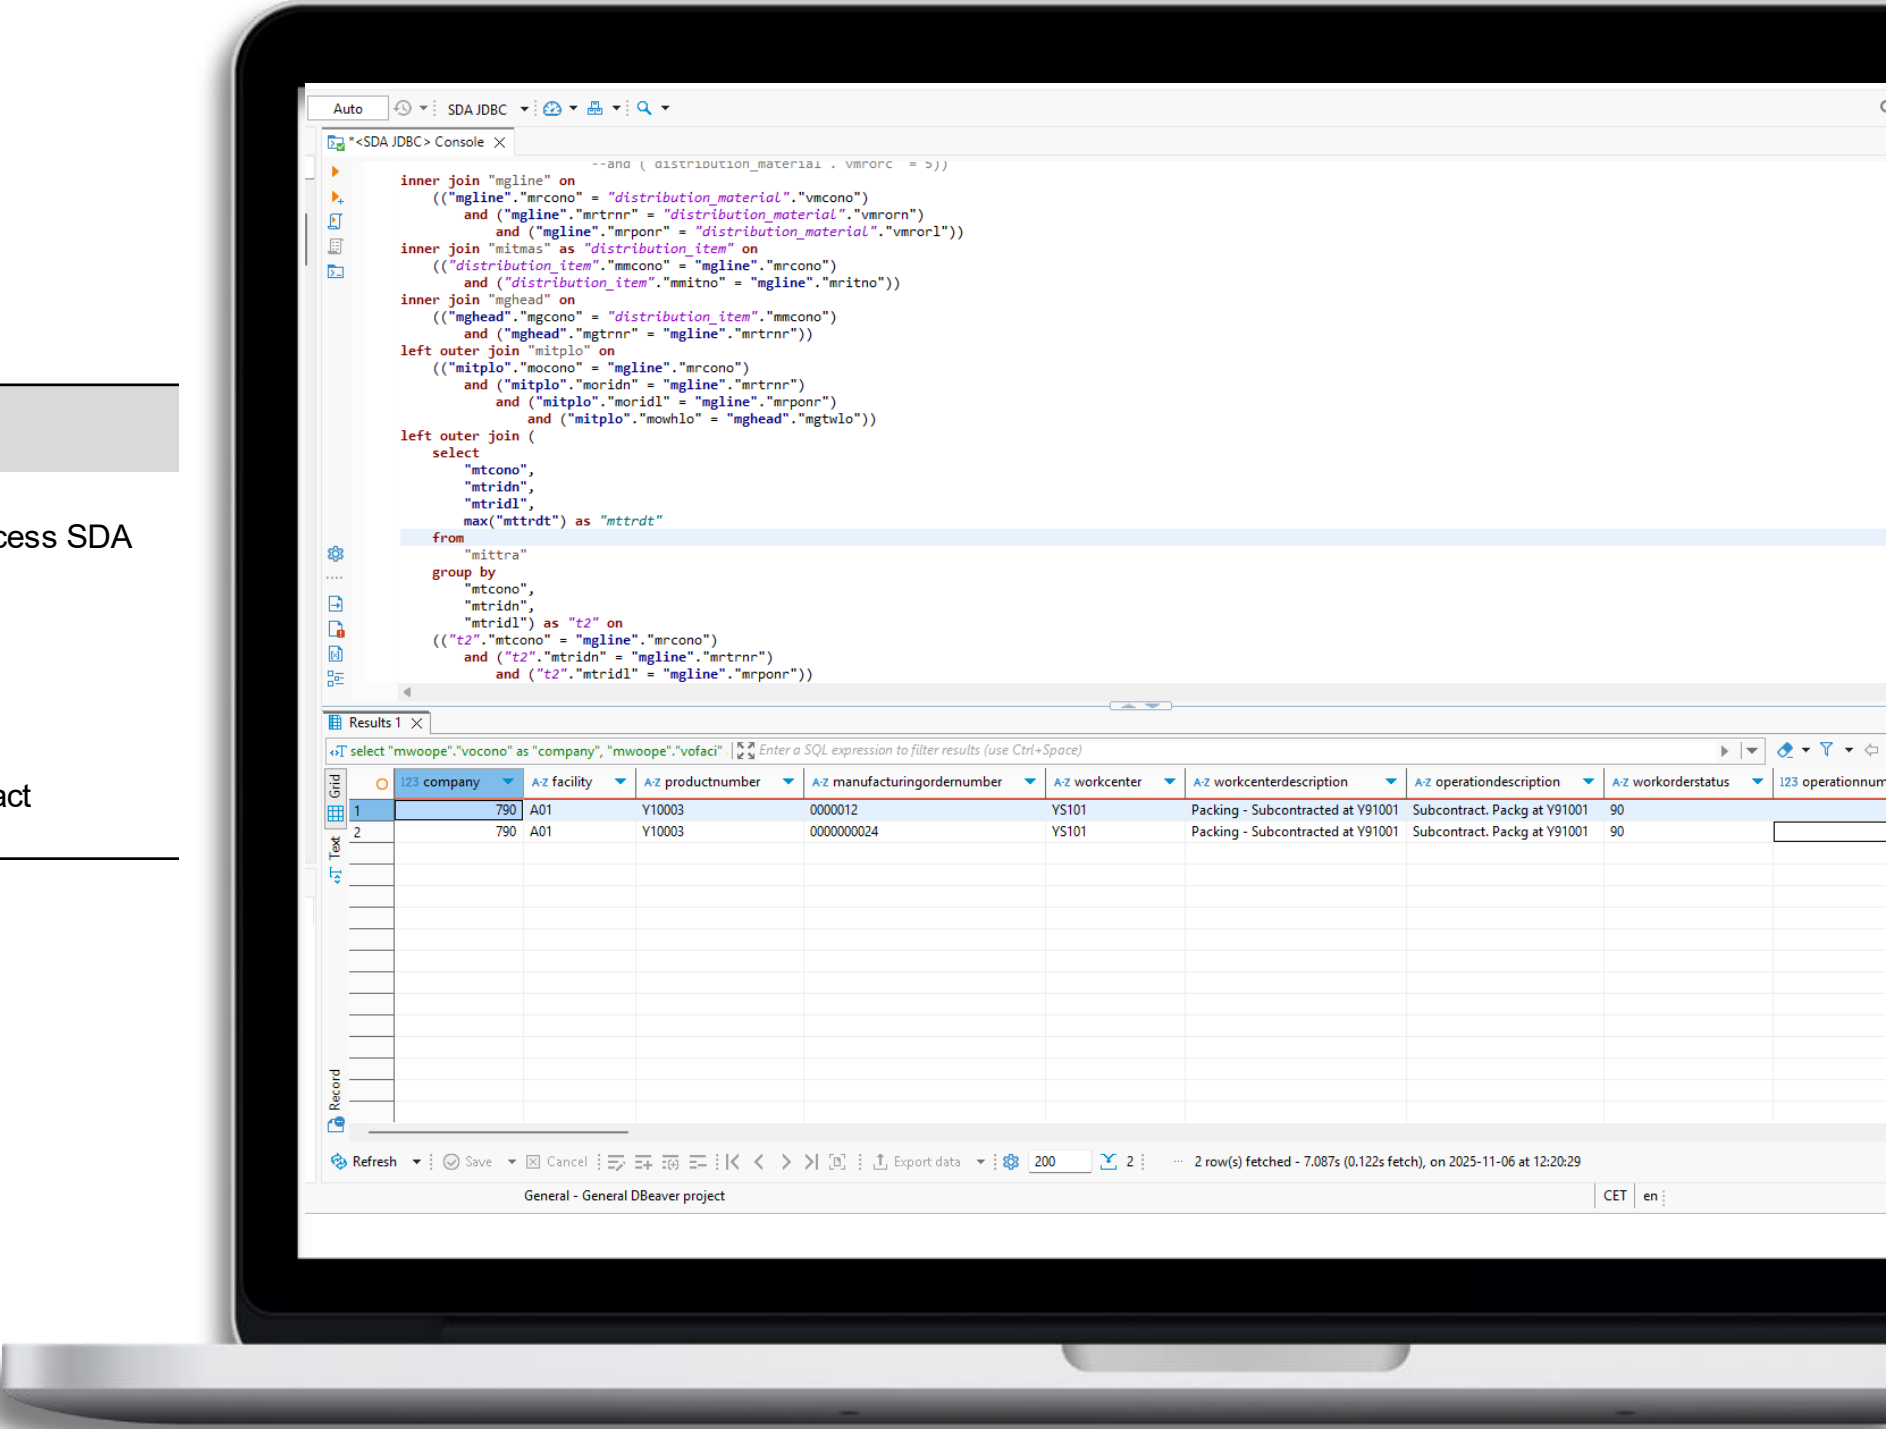
Read Node
Real-time
Safe read-only access
Guardrails for protection

Compass (Infor OS)
Secure access point
Authentication & auditing

Query Tool
SQL interface

How SDA Controls Query Access

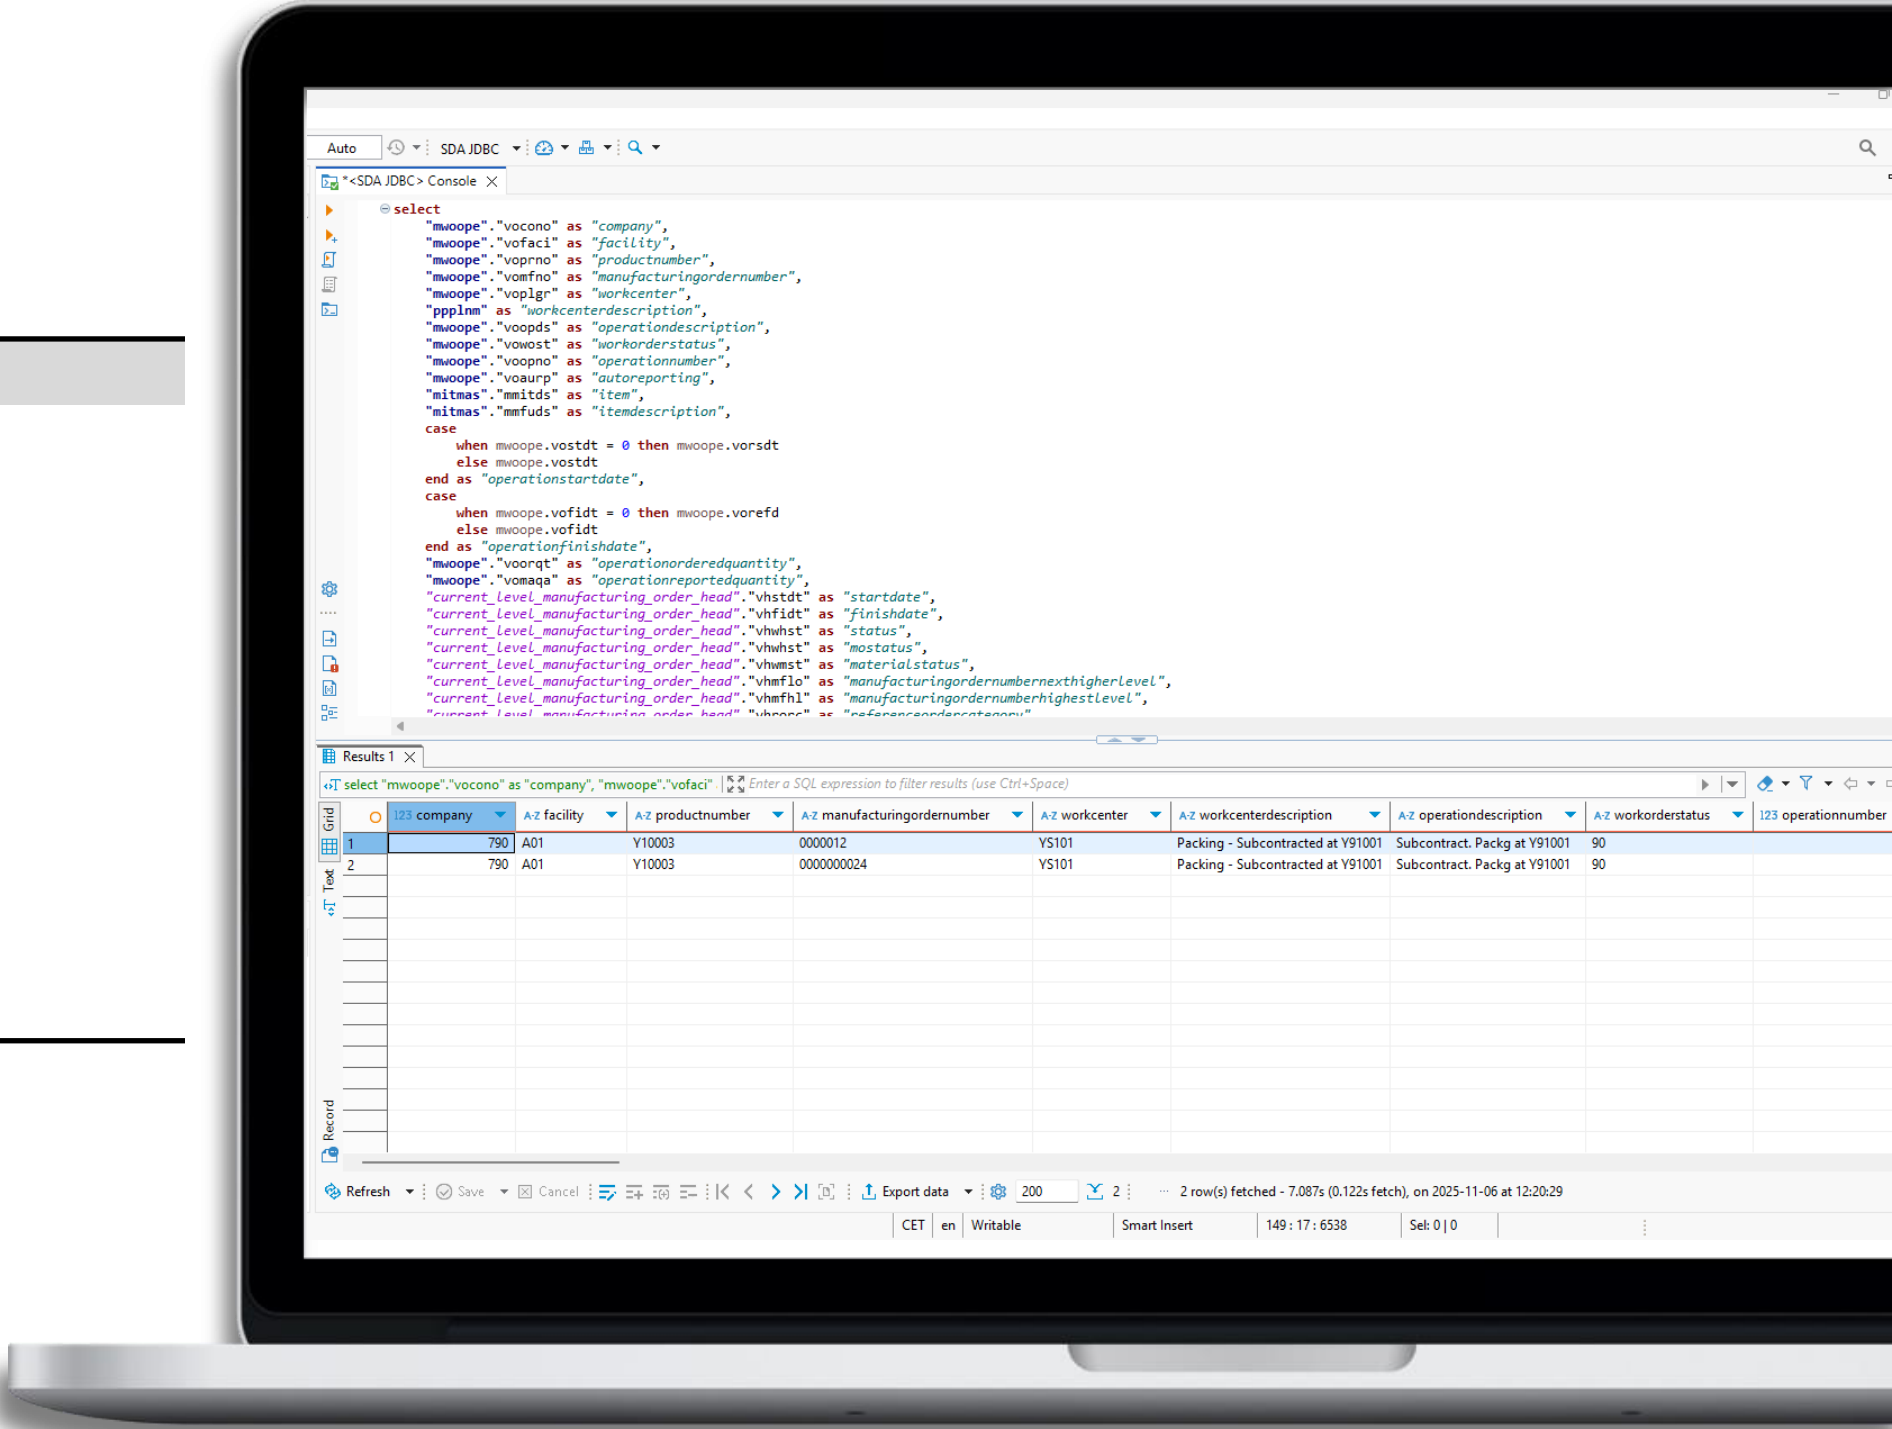
Stage	What happens	What it ensures
Authenticate	User signs in via Infor OS / Compass	Only approved users can access SDA
Authorize	Roles and policies define what the user can query	Controlled access aligned to governance
Execute (Read-only Node)	SQL runs on the read node (not production)	No production workload impact



SDA provides governed, real-time visibility to live M3 data - empowering faster validation, reconciliation, and confidence in the cloud.

Guardrails and Safe Use

Guardrail	Customer benefit
Row limit (e.g., 50,000 rows)	Prevents heavy queries from impacting the service
Query timeout (e.g., 2 minutes)	Predictable performance and fast diagnostics
Concurrency limit (e.g., 10 queries)	Prevents excessive parallel load
Read-only isolation	Queries never run on the live M3 production database
Authentication & audit logging	Secure, traceable, and compliant access



Example Use Cases



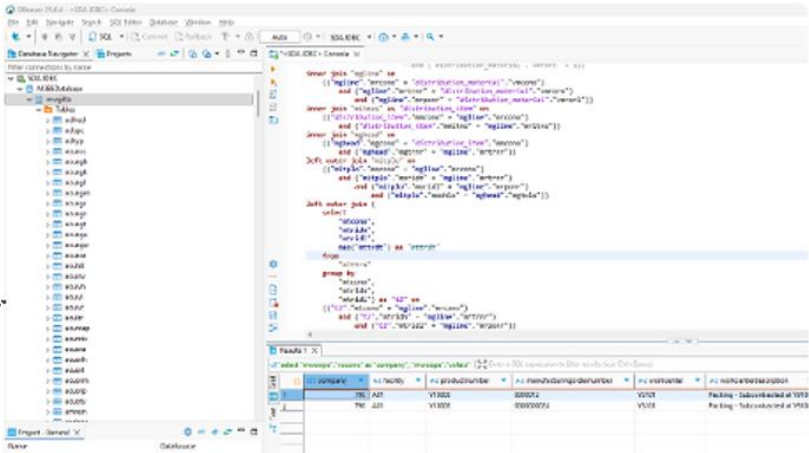
M3



Source Data Access (SDA)

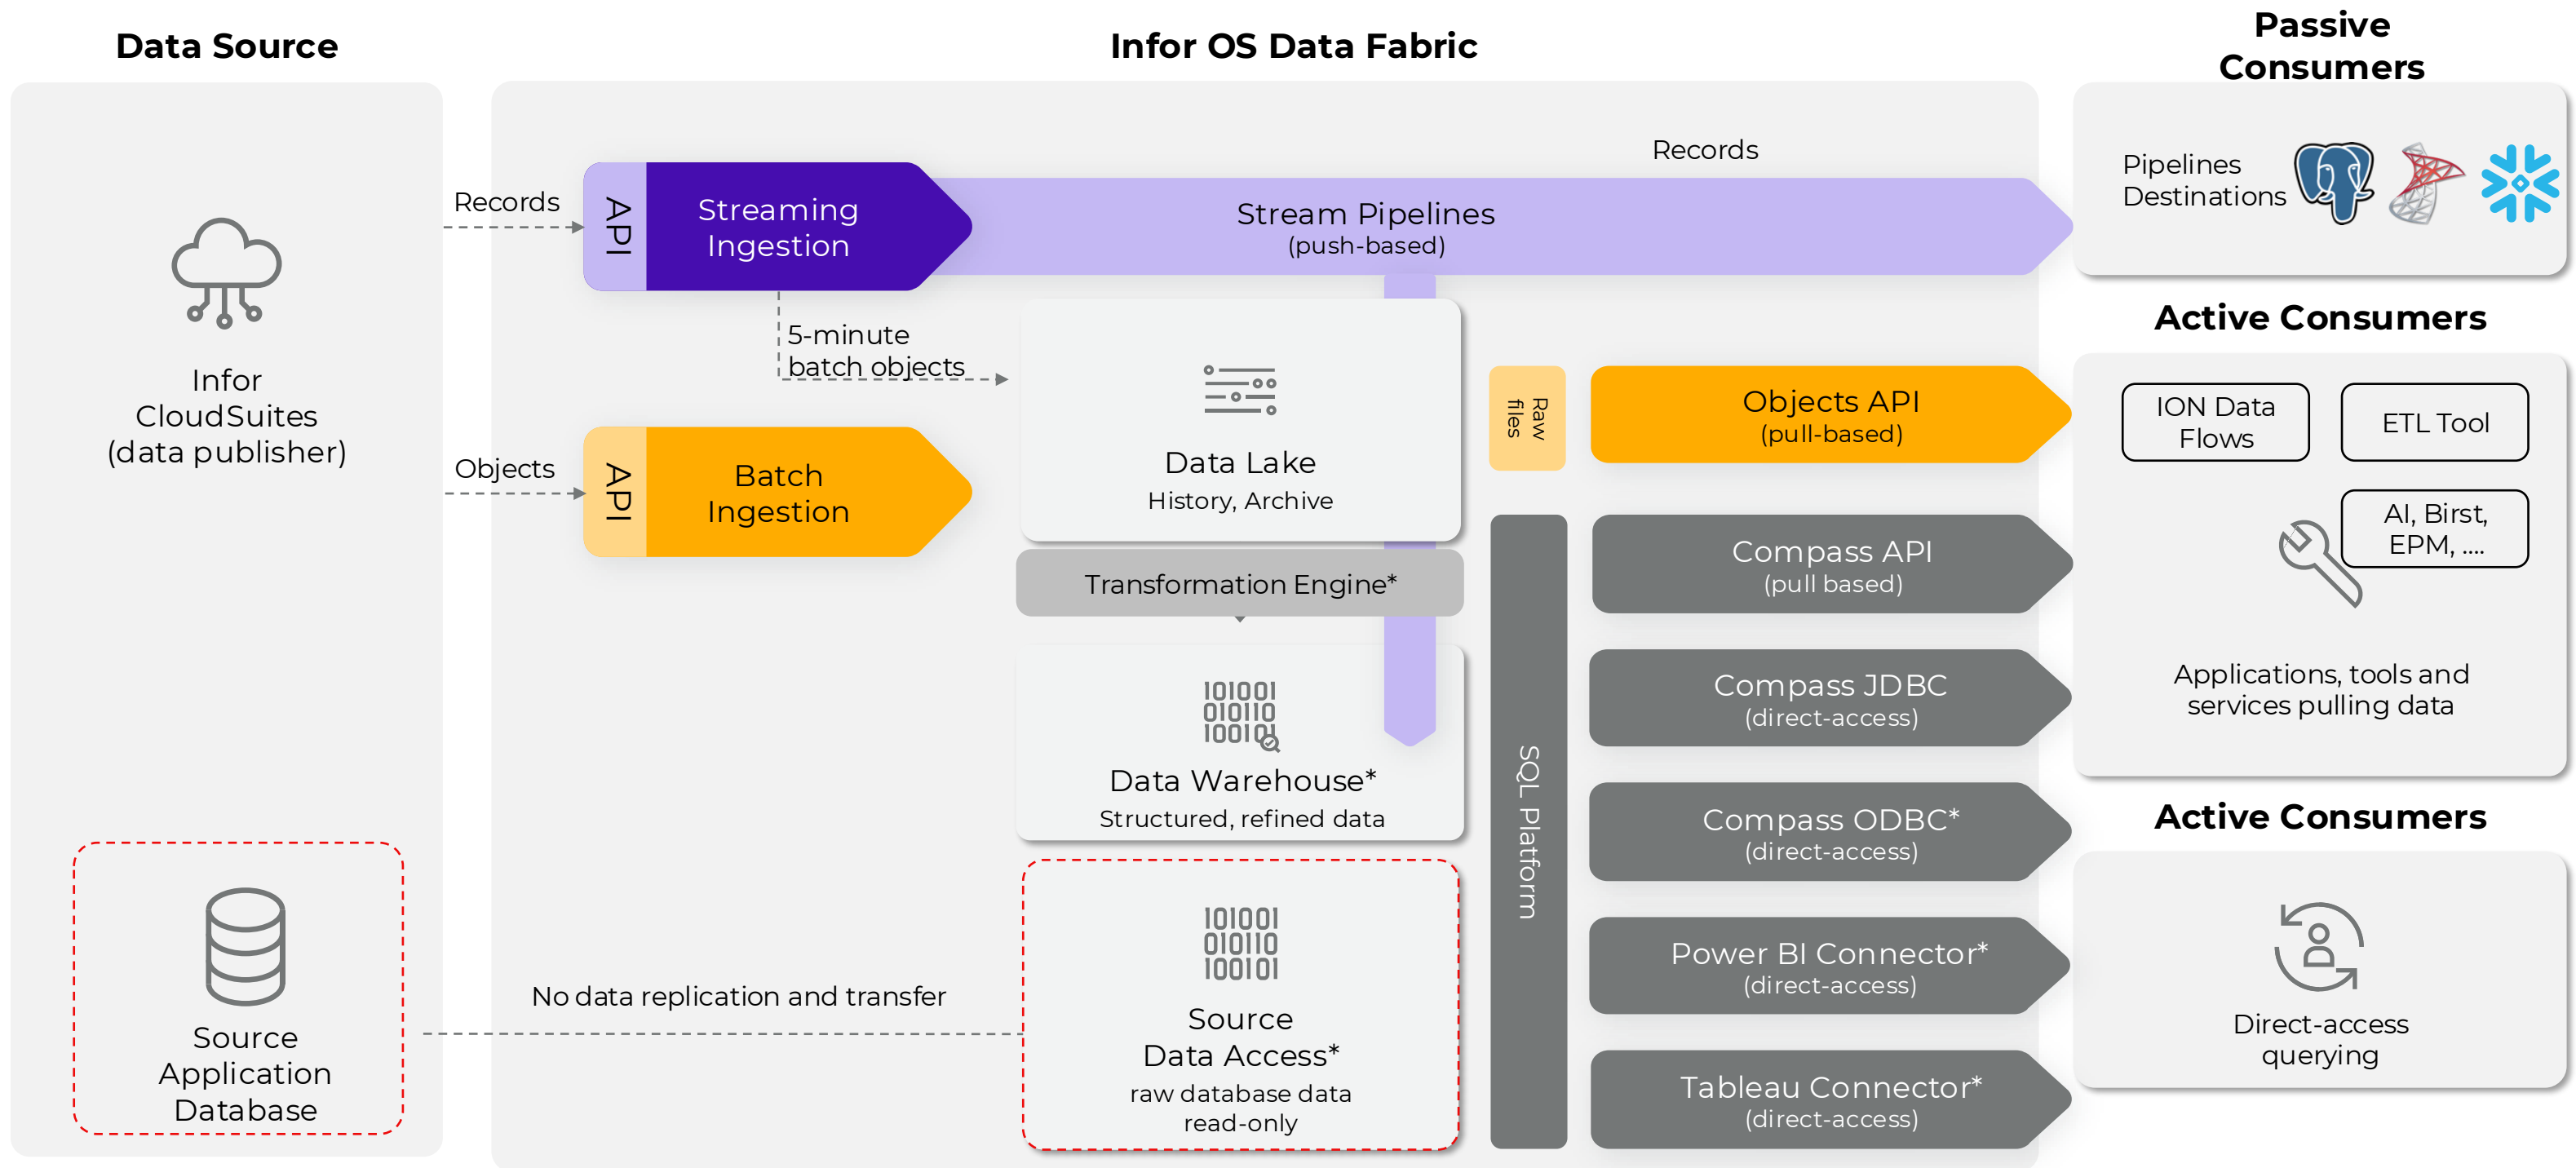


Compass (secure gateway)



Use case	Example questions customers ask	Value / Benefit
Ad hoc investigation	“I need to confirm the data I entered in M3.”	Faster variance analysis
Post-migration validation	“Has all master data migrated correctly?”	Reduced migration risk
Data reconciliation	“Do these figures match what I see in BI?”	Trusted financial reporting
Configuration & process checks	“Show me transactions with the new order type.”	Fewer process exceptions
Audit & compliance	“Do these values add up in the General Ledger?”	Stronger financial controls

Data Access Overview



* Roadmap features

infor[®]