

M3 and Data Fabric

Streaming – Batch loading Data Lake – Metagraph

Infoteam - Den stora teknikdagen 2024

Per Melander
Principal Architect, M3 Development

2024-01-25



Copyright © 2023. Infor. All Rights Reserved. [infor.com](https://www.infor.com)

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.








M3 and Data Fabric

Realtime access to M3 data through streaming and Stream Pipelines in Data Fabric

Artificial Intelligence and Analytics on top of Data Lake as **Big Data** platform

Lakehouse is Infor's **Data Warehouse** as a Service

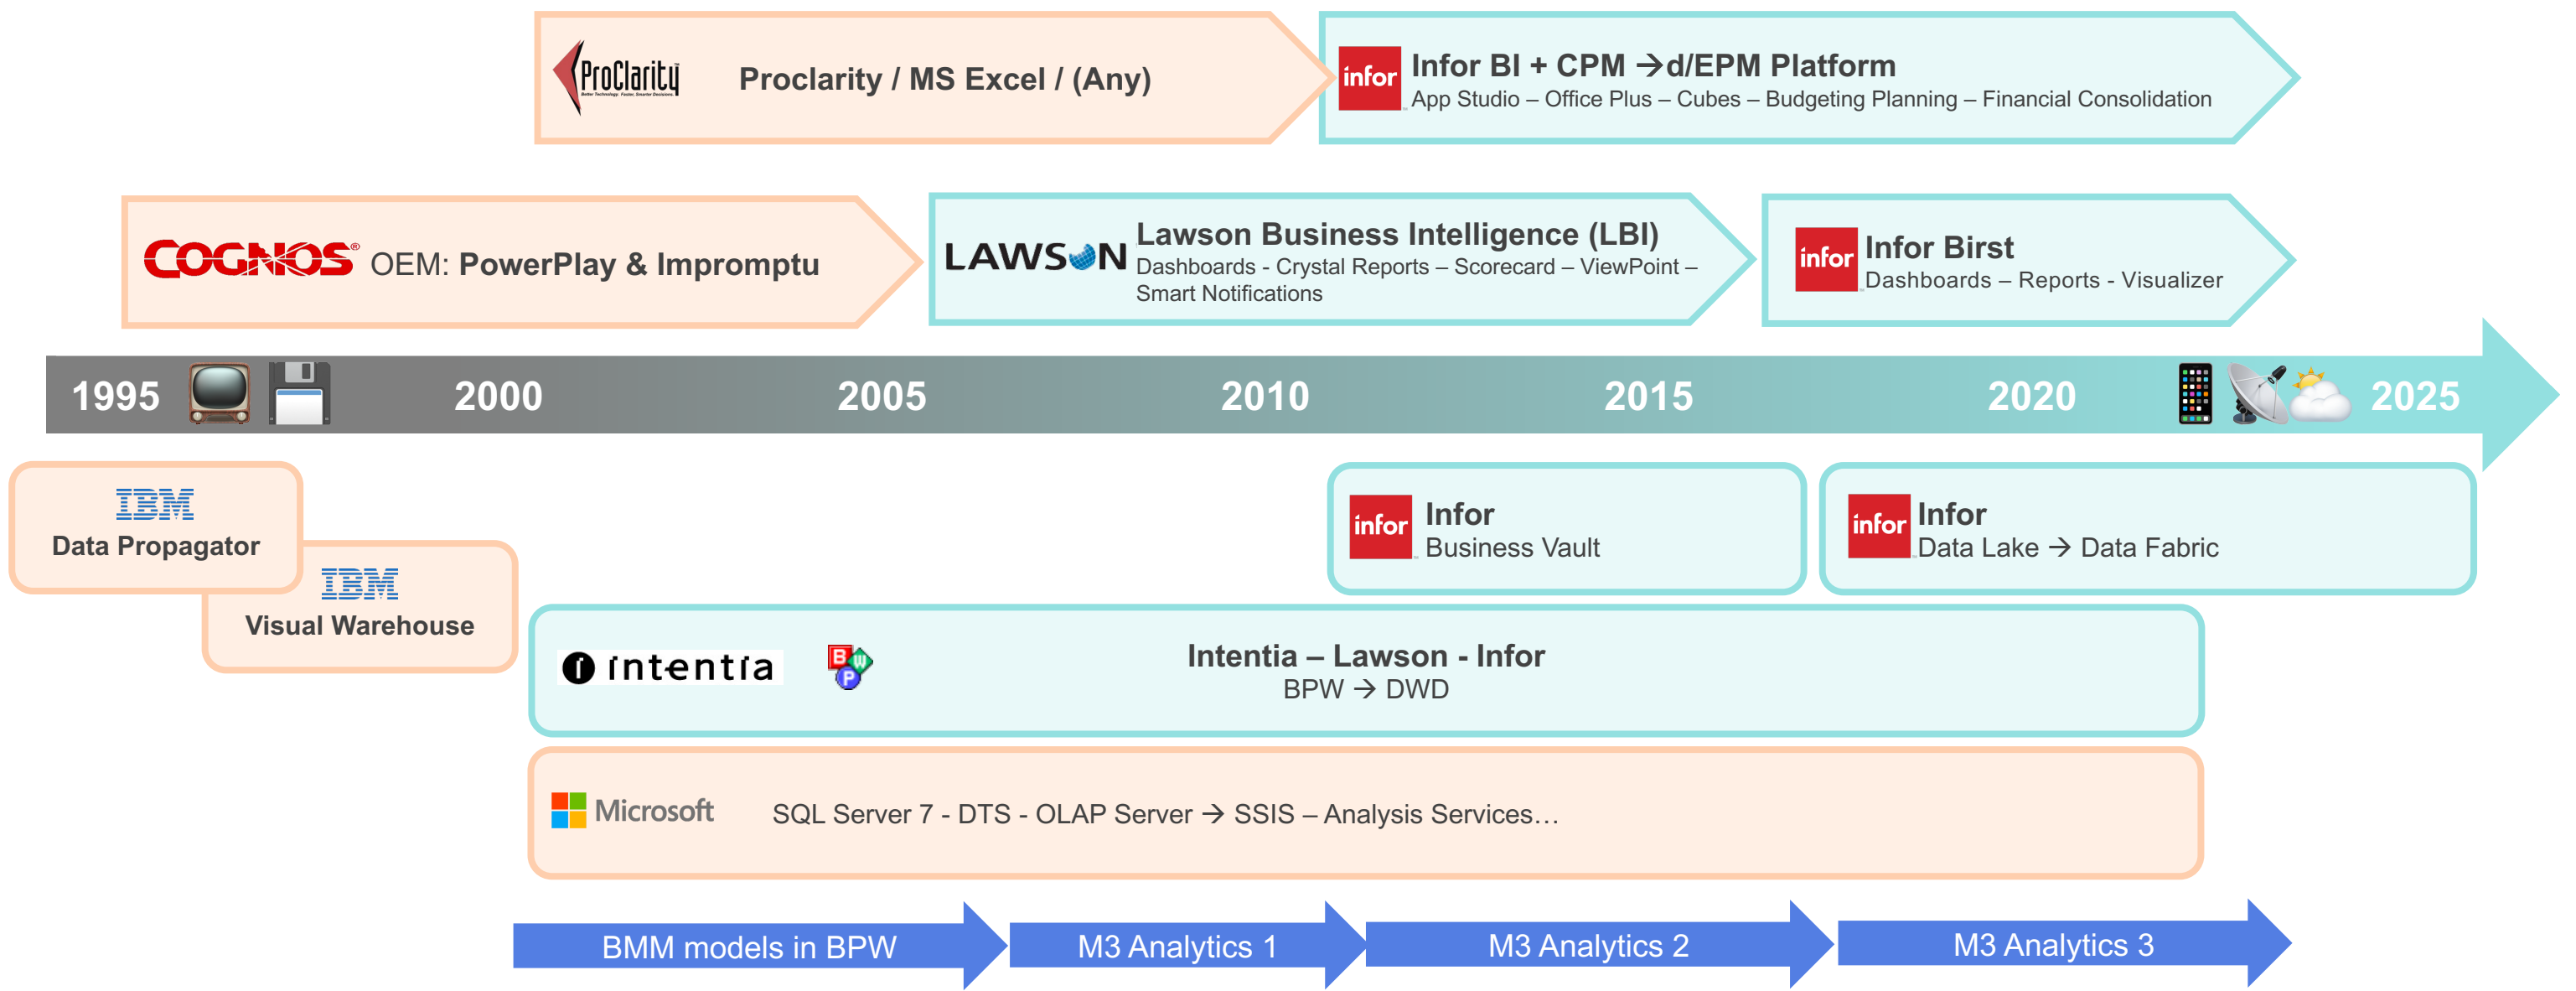
Agenda



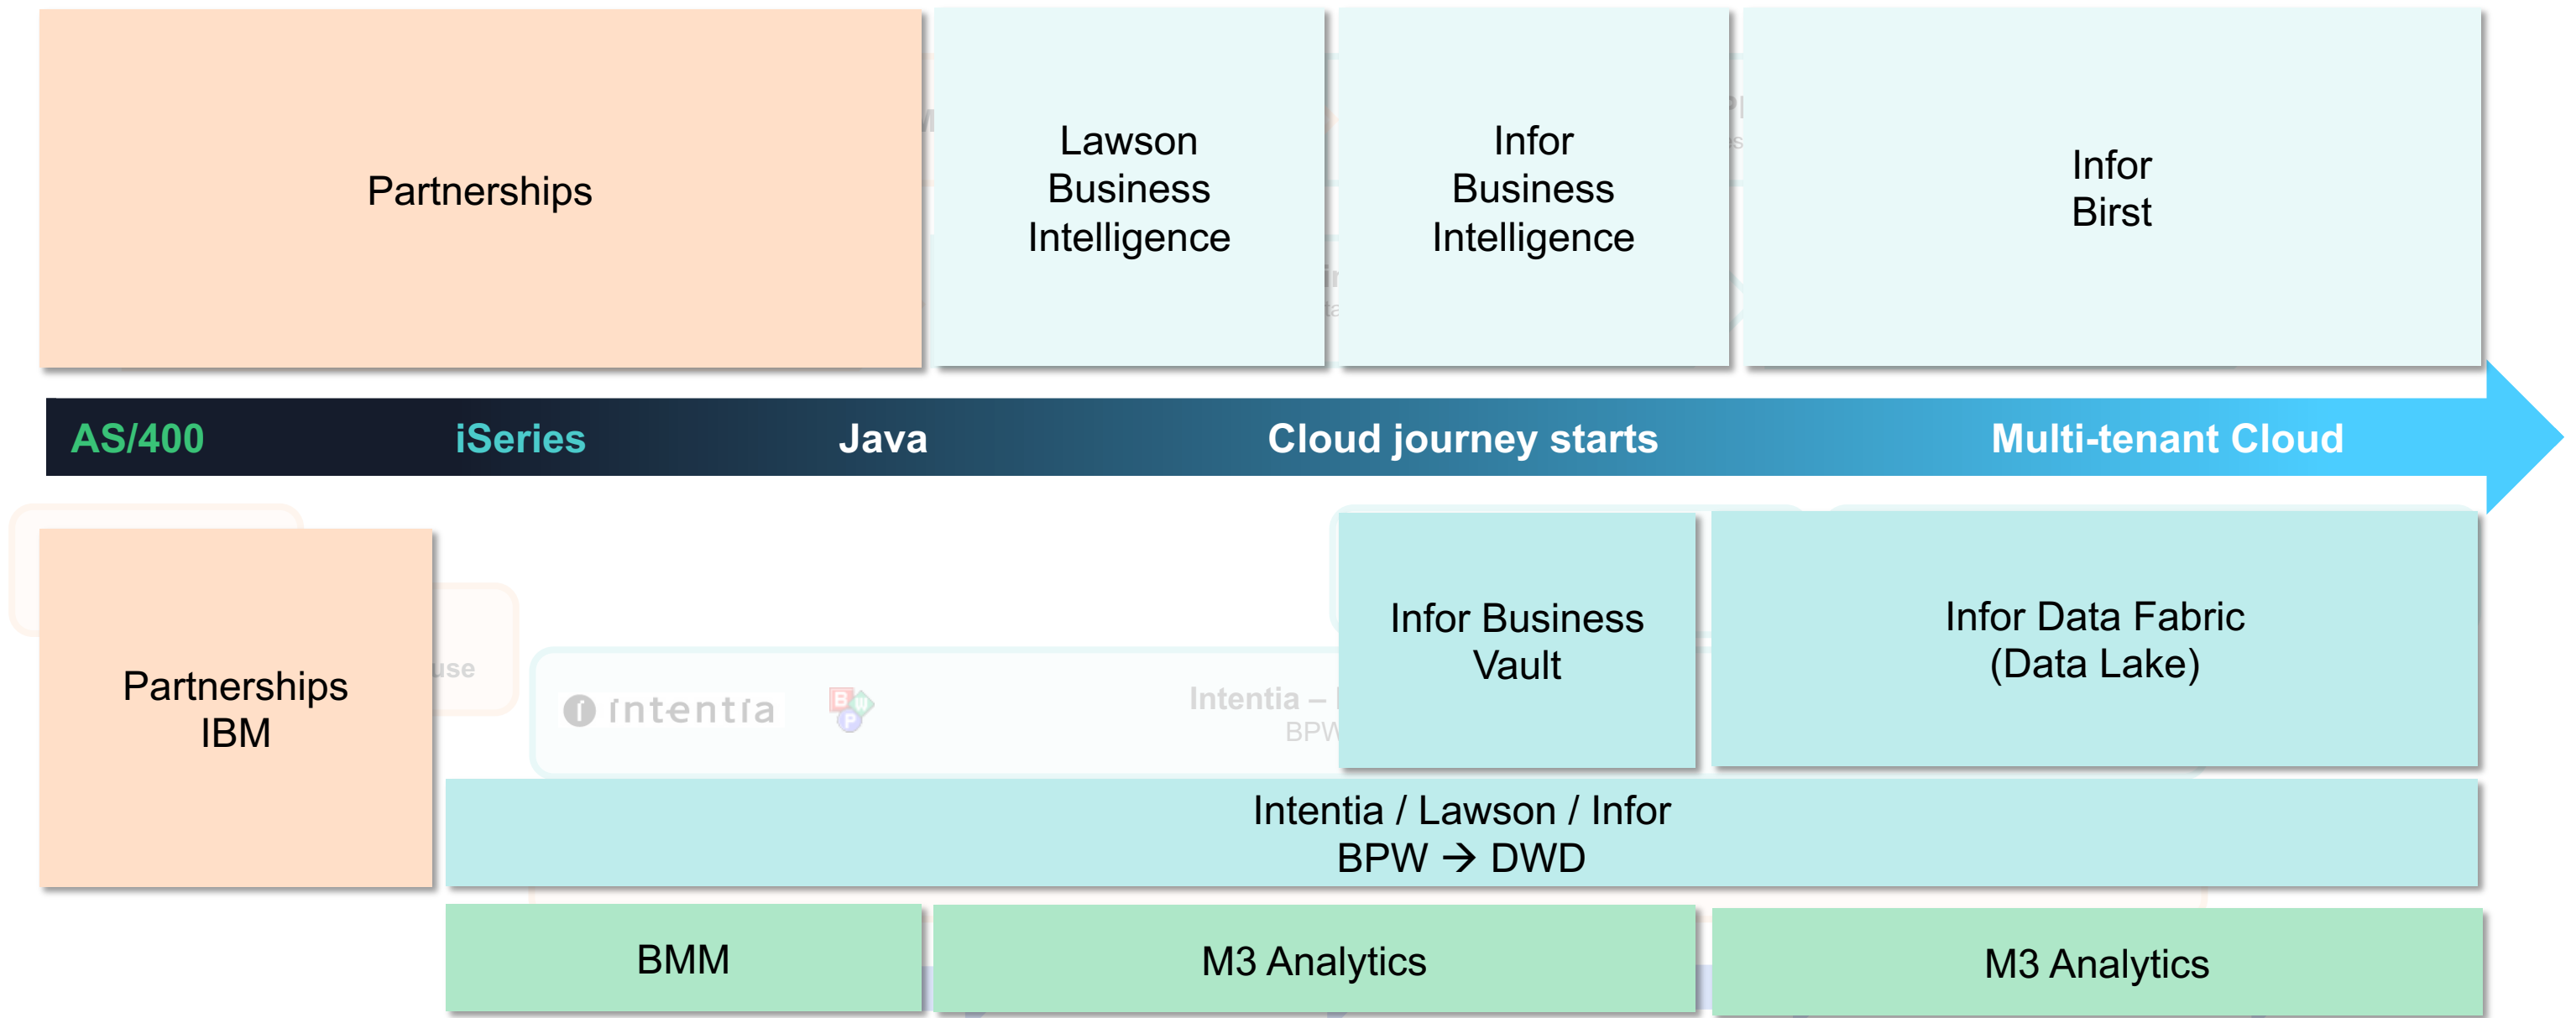
01 Background



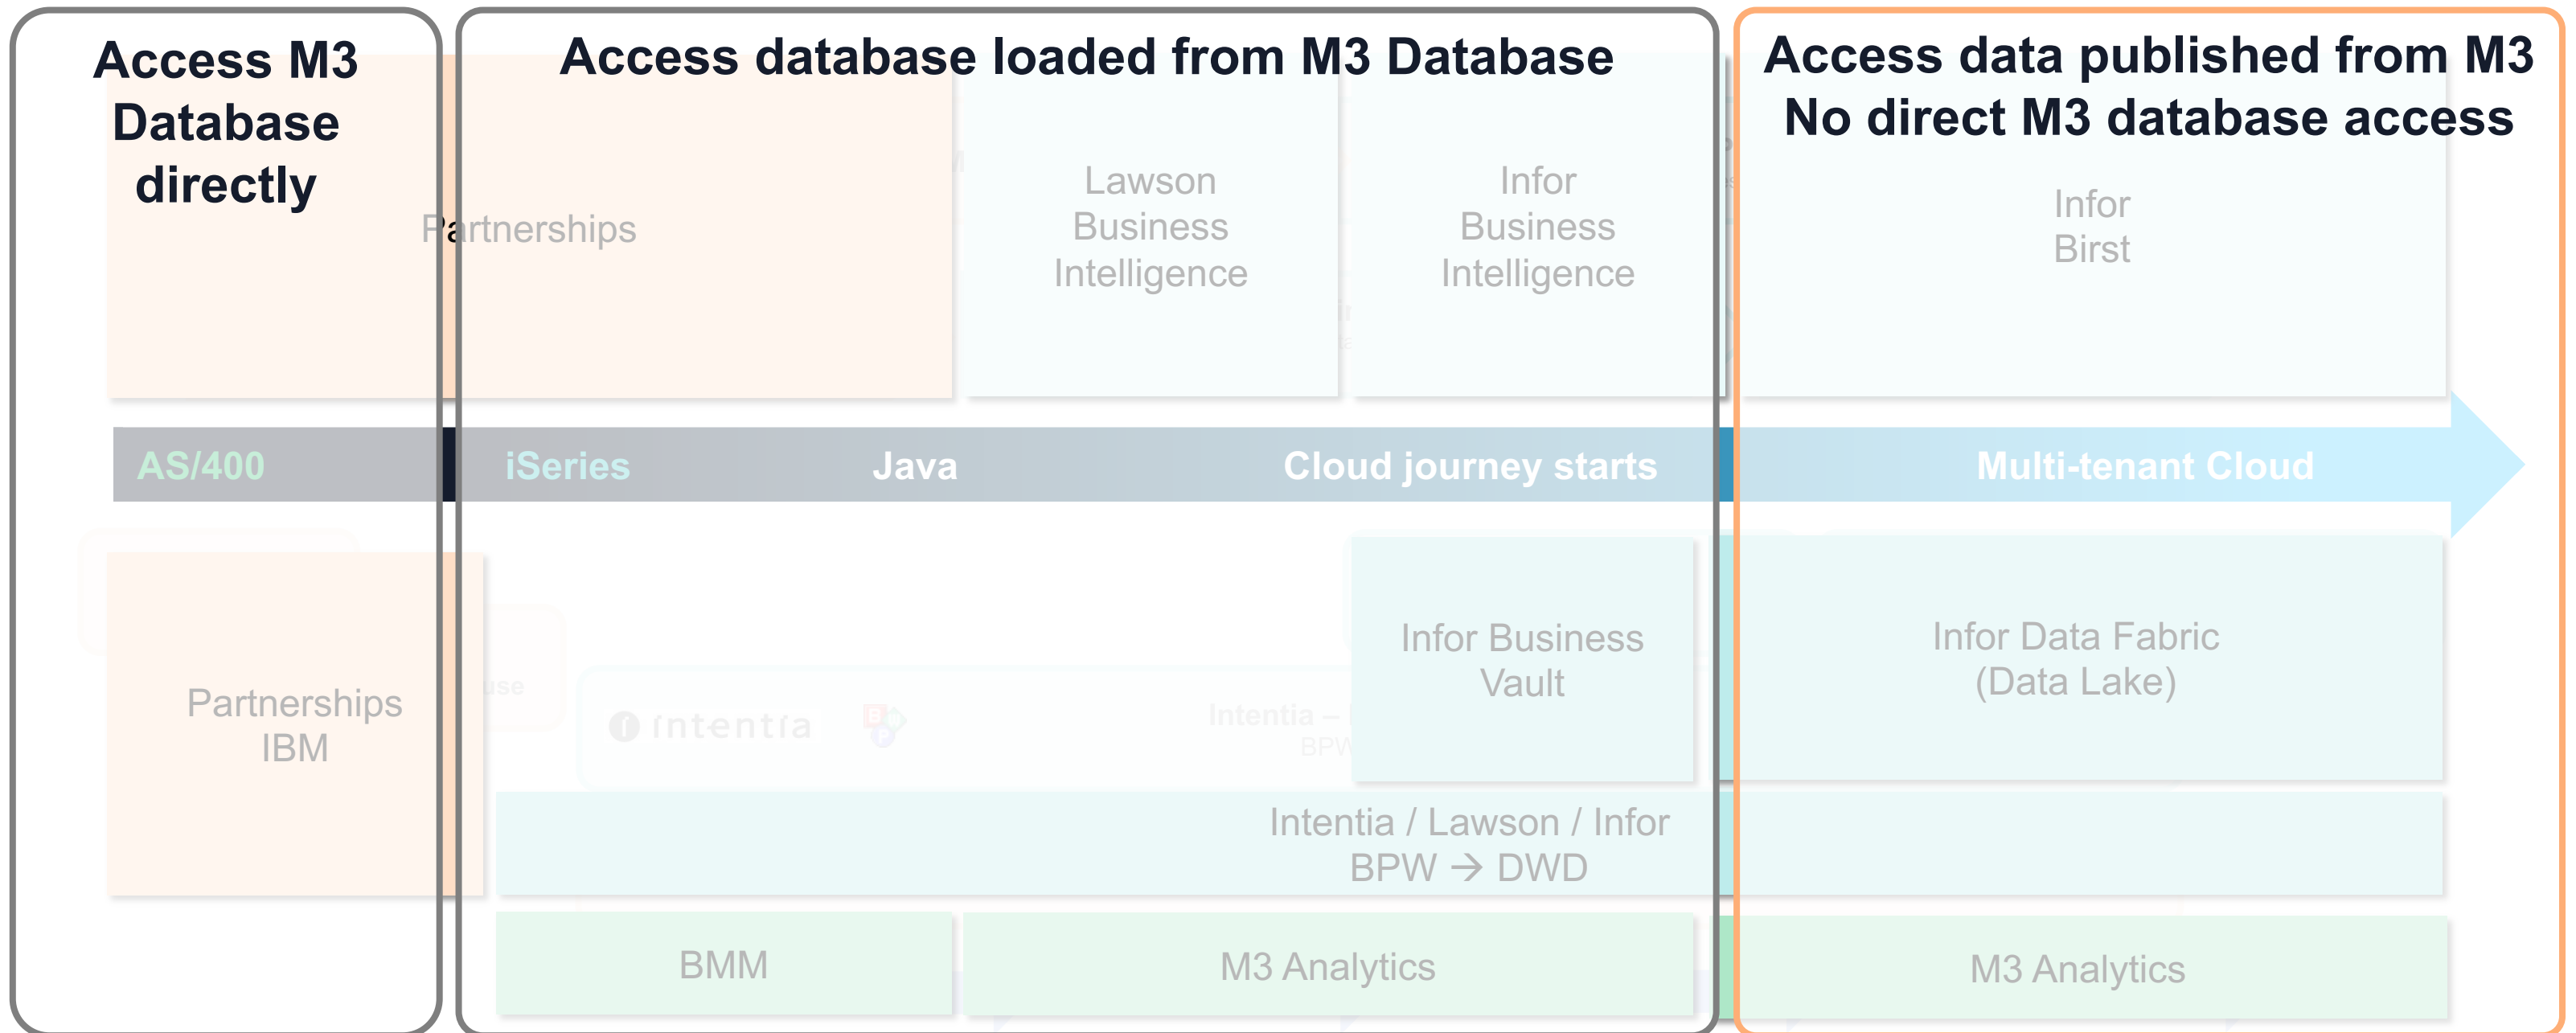
Analytics – BI and Data Warehousing platforms M3 Evolution



Analytics – BI and Data Warehousing platforms M3 Evolution



Analytics – BI and Data Warehousing platforms M3 Evolution



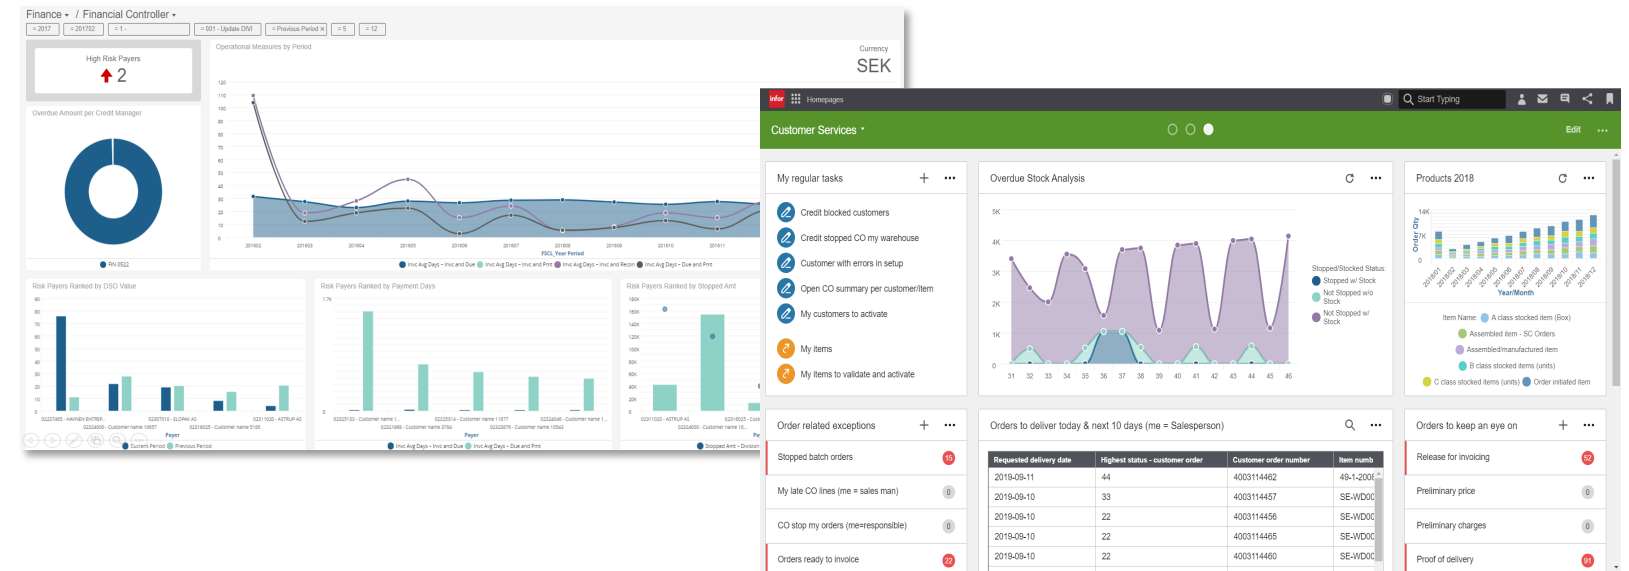
02

What is Data Fabric



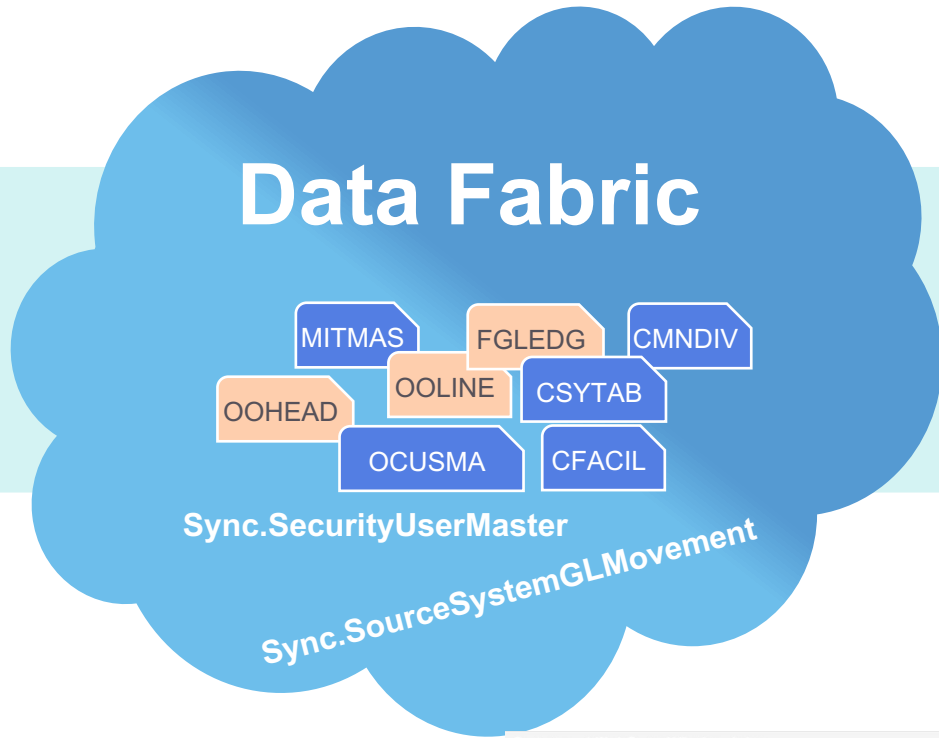
M3 and Data Fabric

Why Data Fabric?



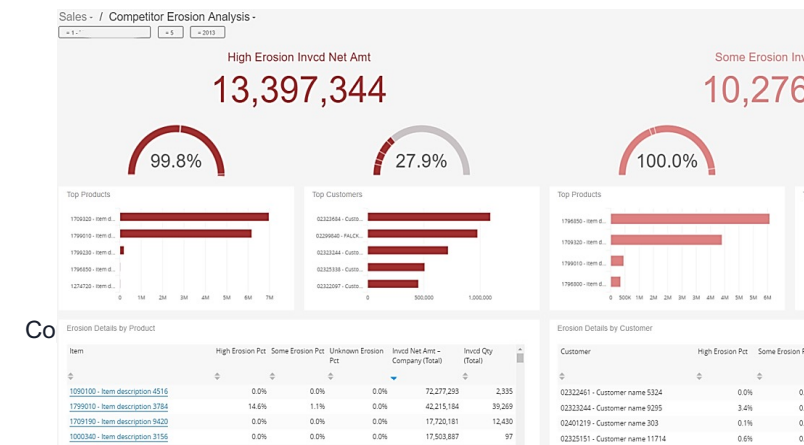
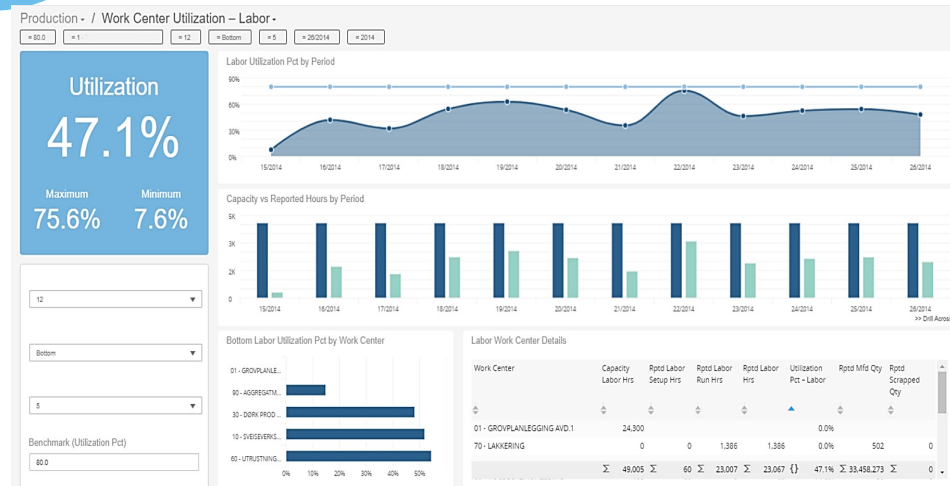

Infor M3
CloudSuite
Applications

Push

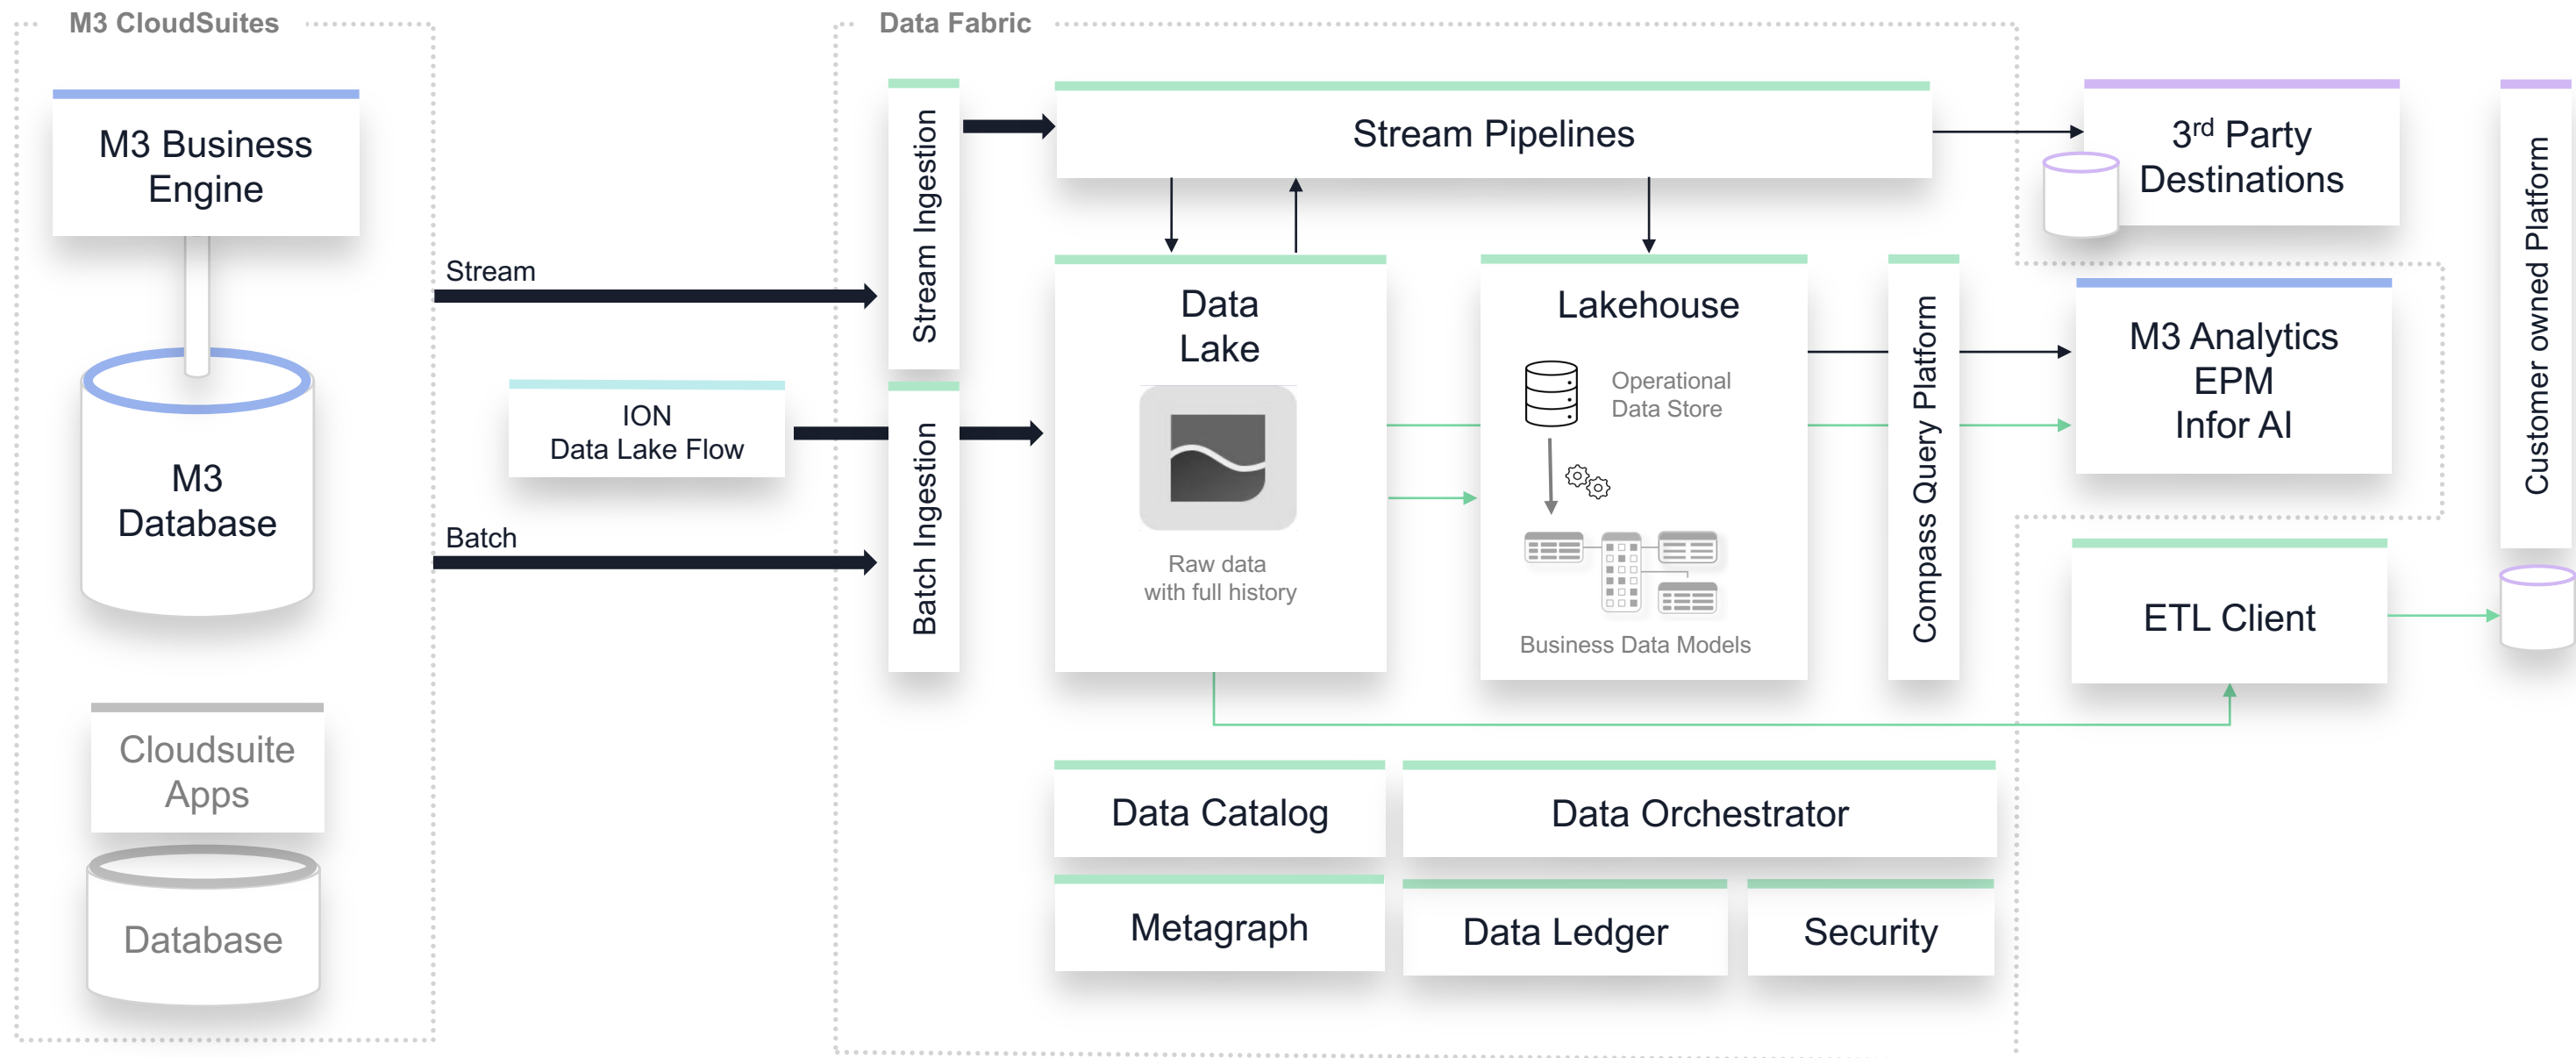


Access

- M3 Analytics
- EPM
- Infor AI
- GRC
- Data Warehouse



M3 and Data Fabric 2024



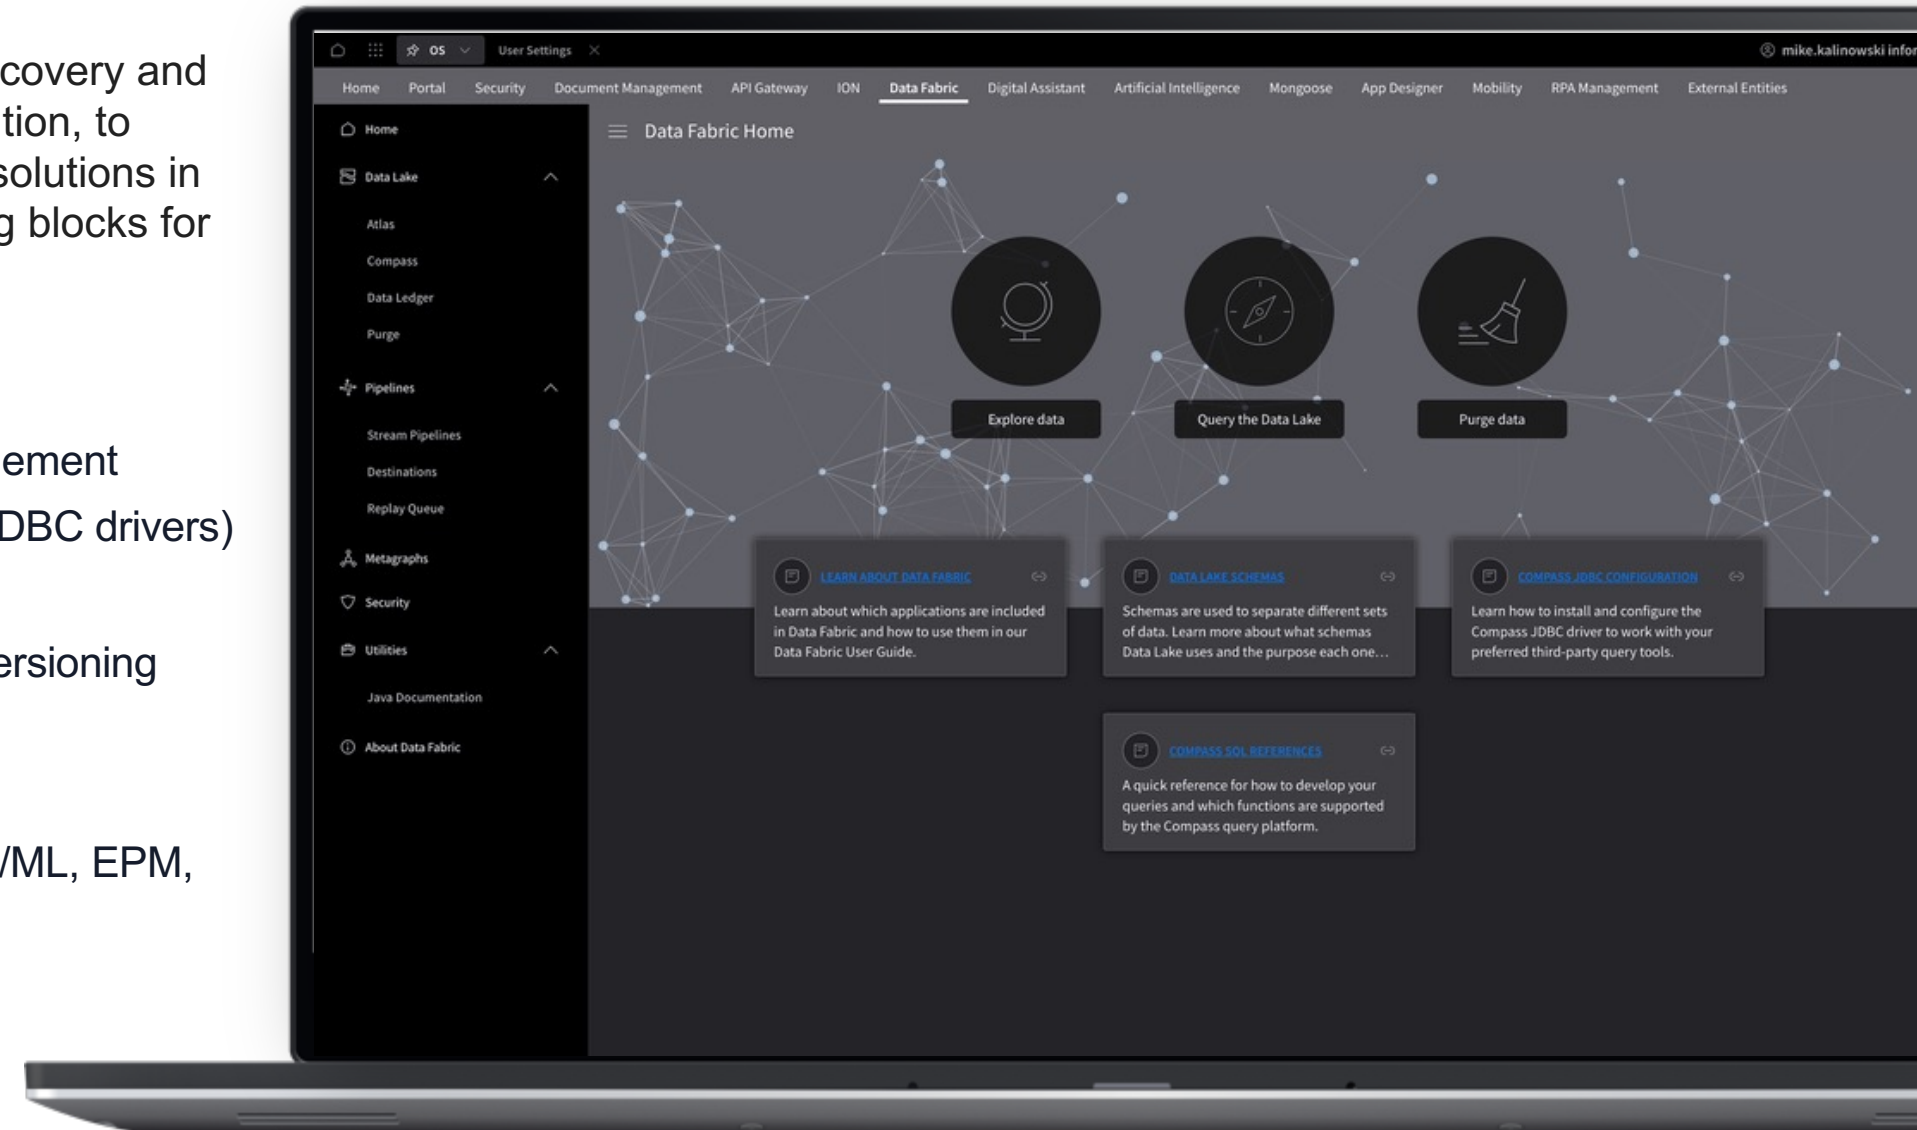
CloudSuite Data Platform

M3 & Data Fabric

M3 will ultimately build an open content framework for efficient discovery and utilization in Lakehouse, Infor's Data warehouse-as-a-service solution, to support both operational data use cases as well as rich analytics solutions in an open and flexible platform providing predefined content building blocks for Analytics, EPM, Coleman AI and 3rd party BI platforms.

Key features

- Simple point & click data publishing and data integration management
- Compass SQL – ANSI SQL query interfaces (RESTful APIs & JDBC drivers)
- Consolidated metadata management in Data Catalog
- Query across time to explore how data has evolved with data versioning
- Real-time and batch-based ingestion platform
- Stream Pipelines for real-time delivery and operational needs
- Seamless data-driven integration ecosystem (Birst, Coleman AI/ML, EPM, Mongoose & more)



M3 Reporting Options



03 Data Fabric Solutions



M3 Solutions using Infor Data Lake

M3 Analytics using the Infor Birst platform

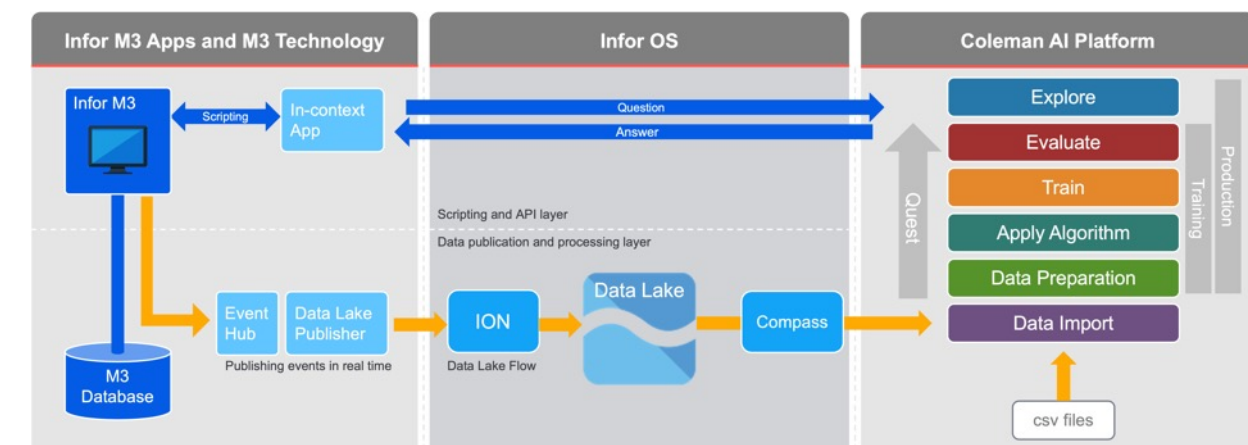
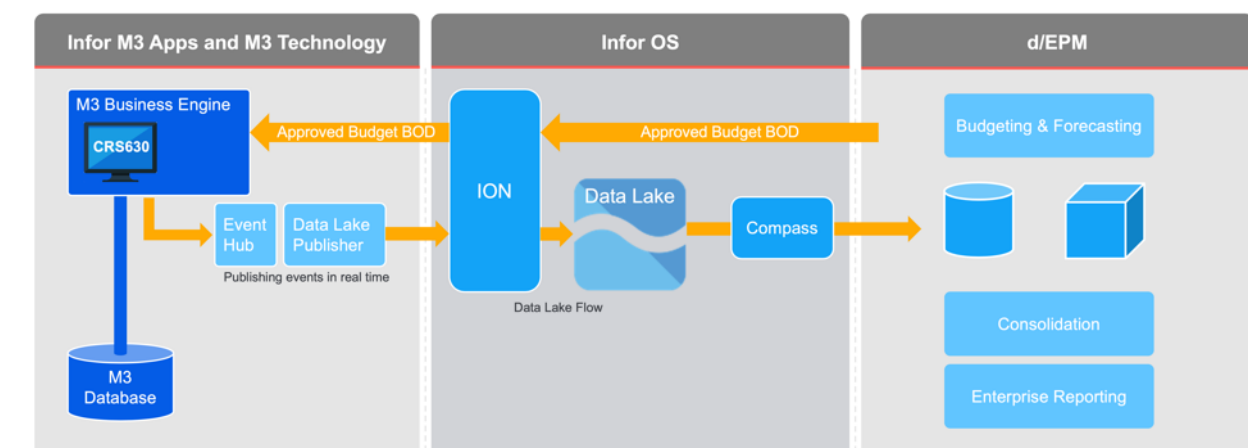
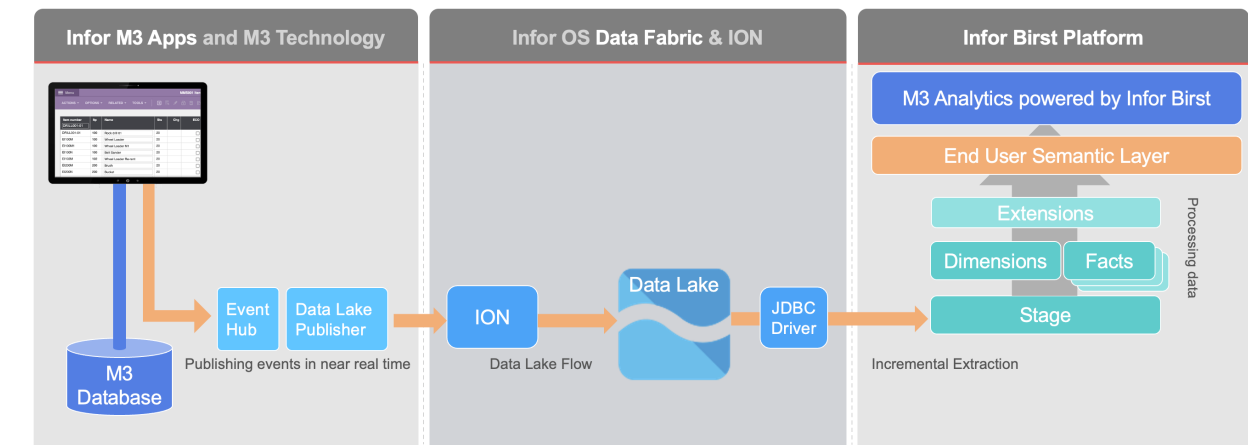
- M3 is feeding Infor Data Lake in near real time – initial load option also available.
- M3 Analytics was the first solution to integrate to Infor Data Lake via Birst.
- First early adopters enabled in August 2019 – General Available in March 2020

M3 CE integration with d/EPM CE

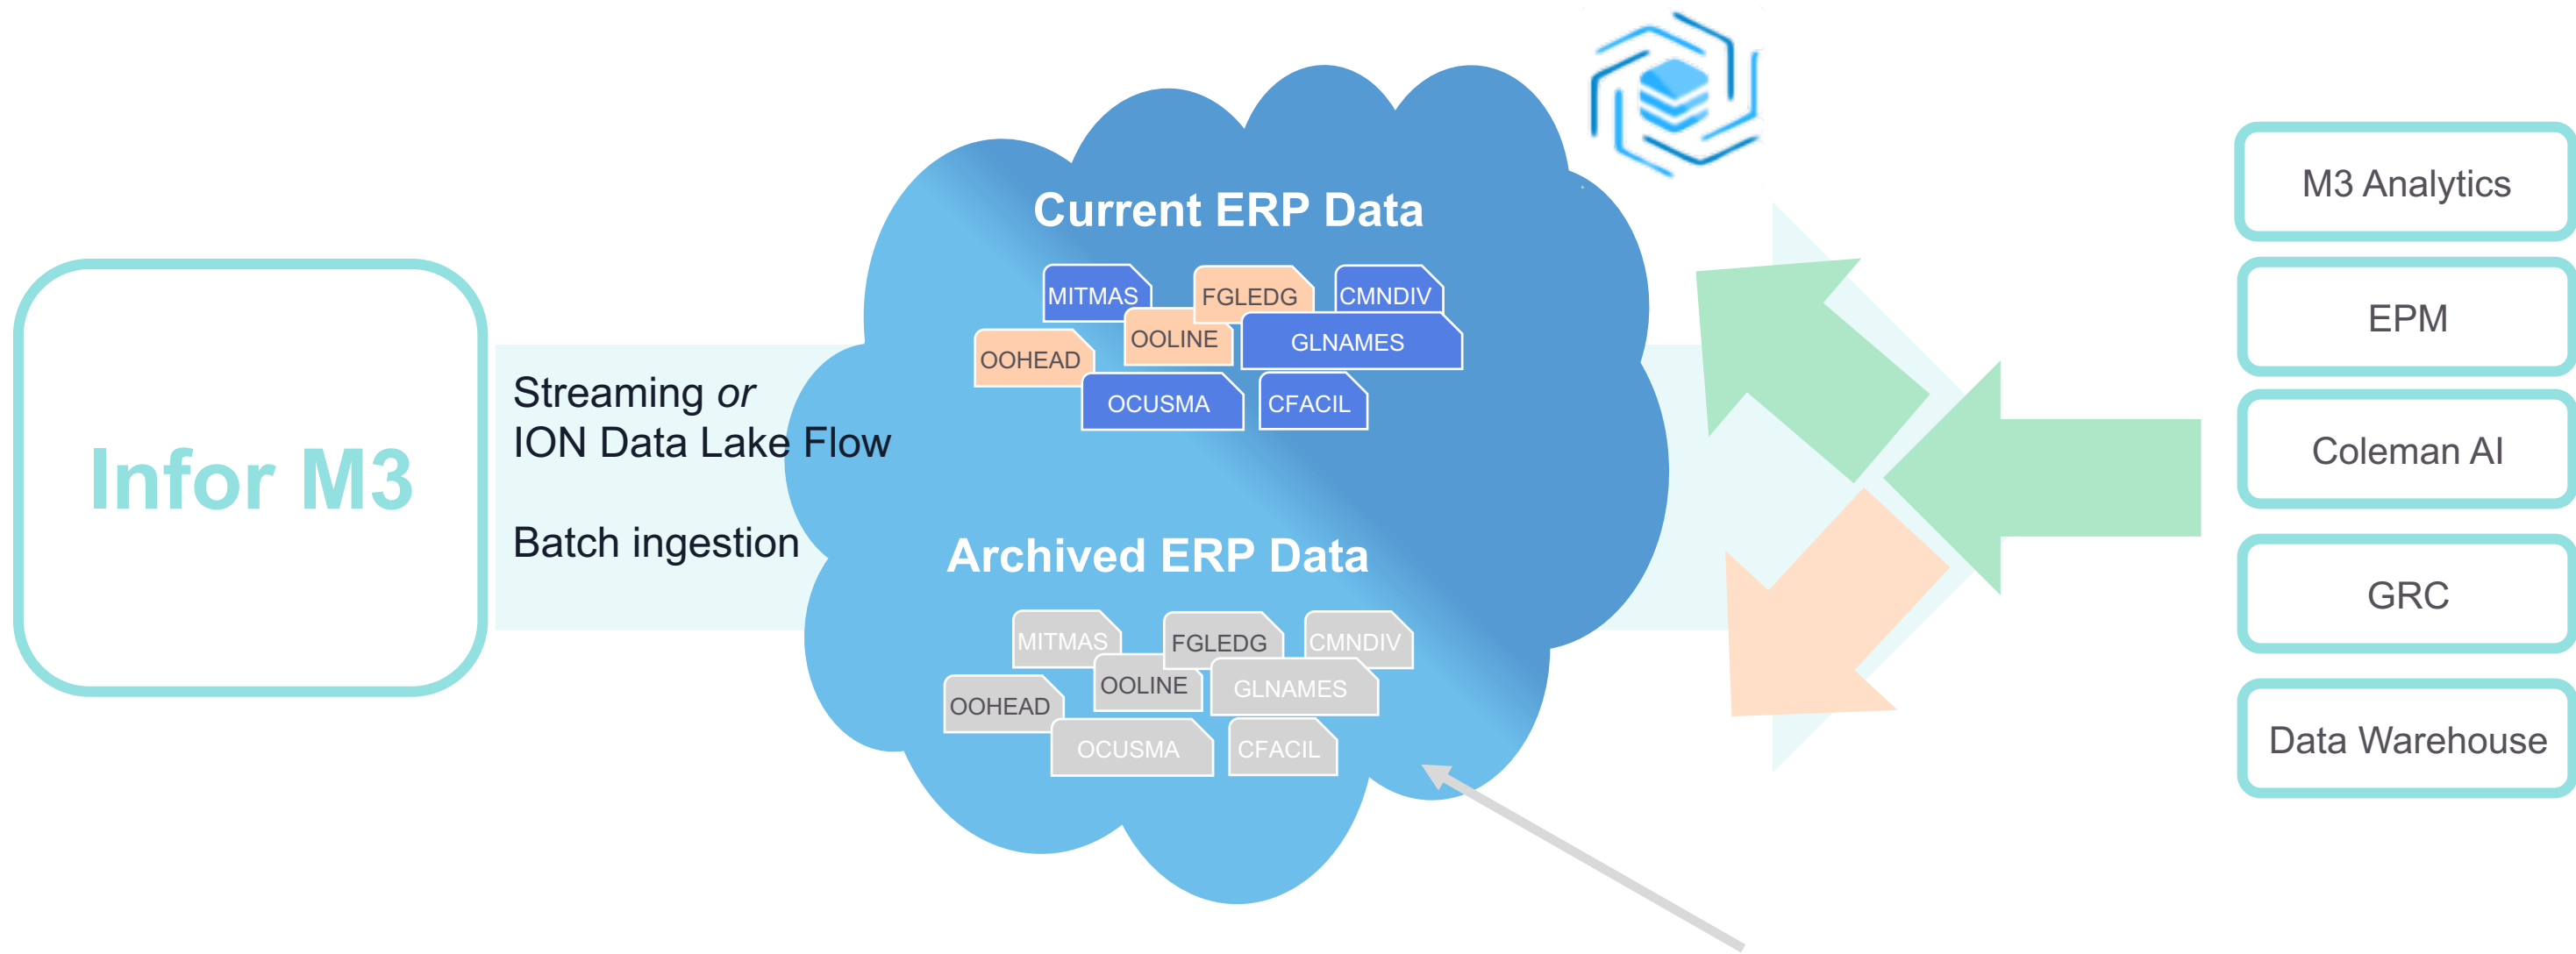
- This two-way integration enables Infor d/EPM to send an approved budget back to M3 to control actual costs and revenues. M3 data is published to Infor Data Lake and Infor d/EPM is loading into staging- and integration tables to feed multidimensional structures.
- This new M3 CE integration was enabled in June 2020
- Financial Consolidation integration was enabled in September 2021

M3 CE integration with Infor AI Platform

- Delivered example: Sales Price Recommender
 - *How much discount should I give this customer for this product if I want to win the order?*
- Training material on Campus
- Delivered via KB Article: Data Lake Queries, Quests, Widget, and M3 IC App instructions
- Available as content in June 2020



Transaction Archiving and Consumption



Transparent access to data from existing solutions by the metadata property ("archive flag") in Data Catalog + Data Lake

M3 Archiving to Data Lake

Offload and save space in your transactional database by moving archived data to Data Lake

Archiving to Data Lake is a functionality in M3 Business Engine that takes archived data into Data Lake.

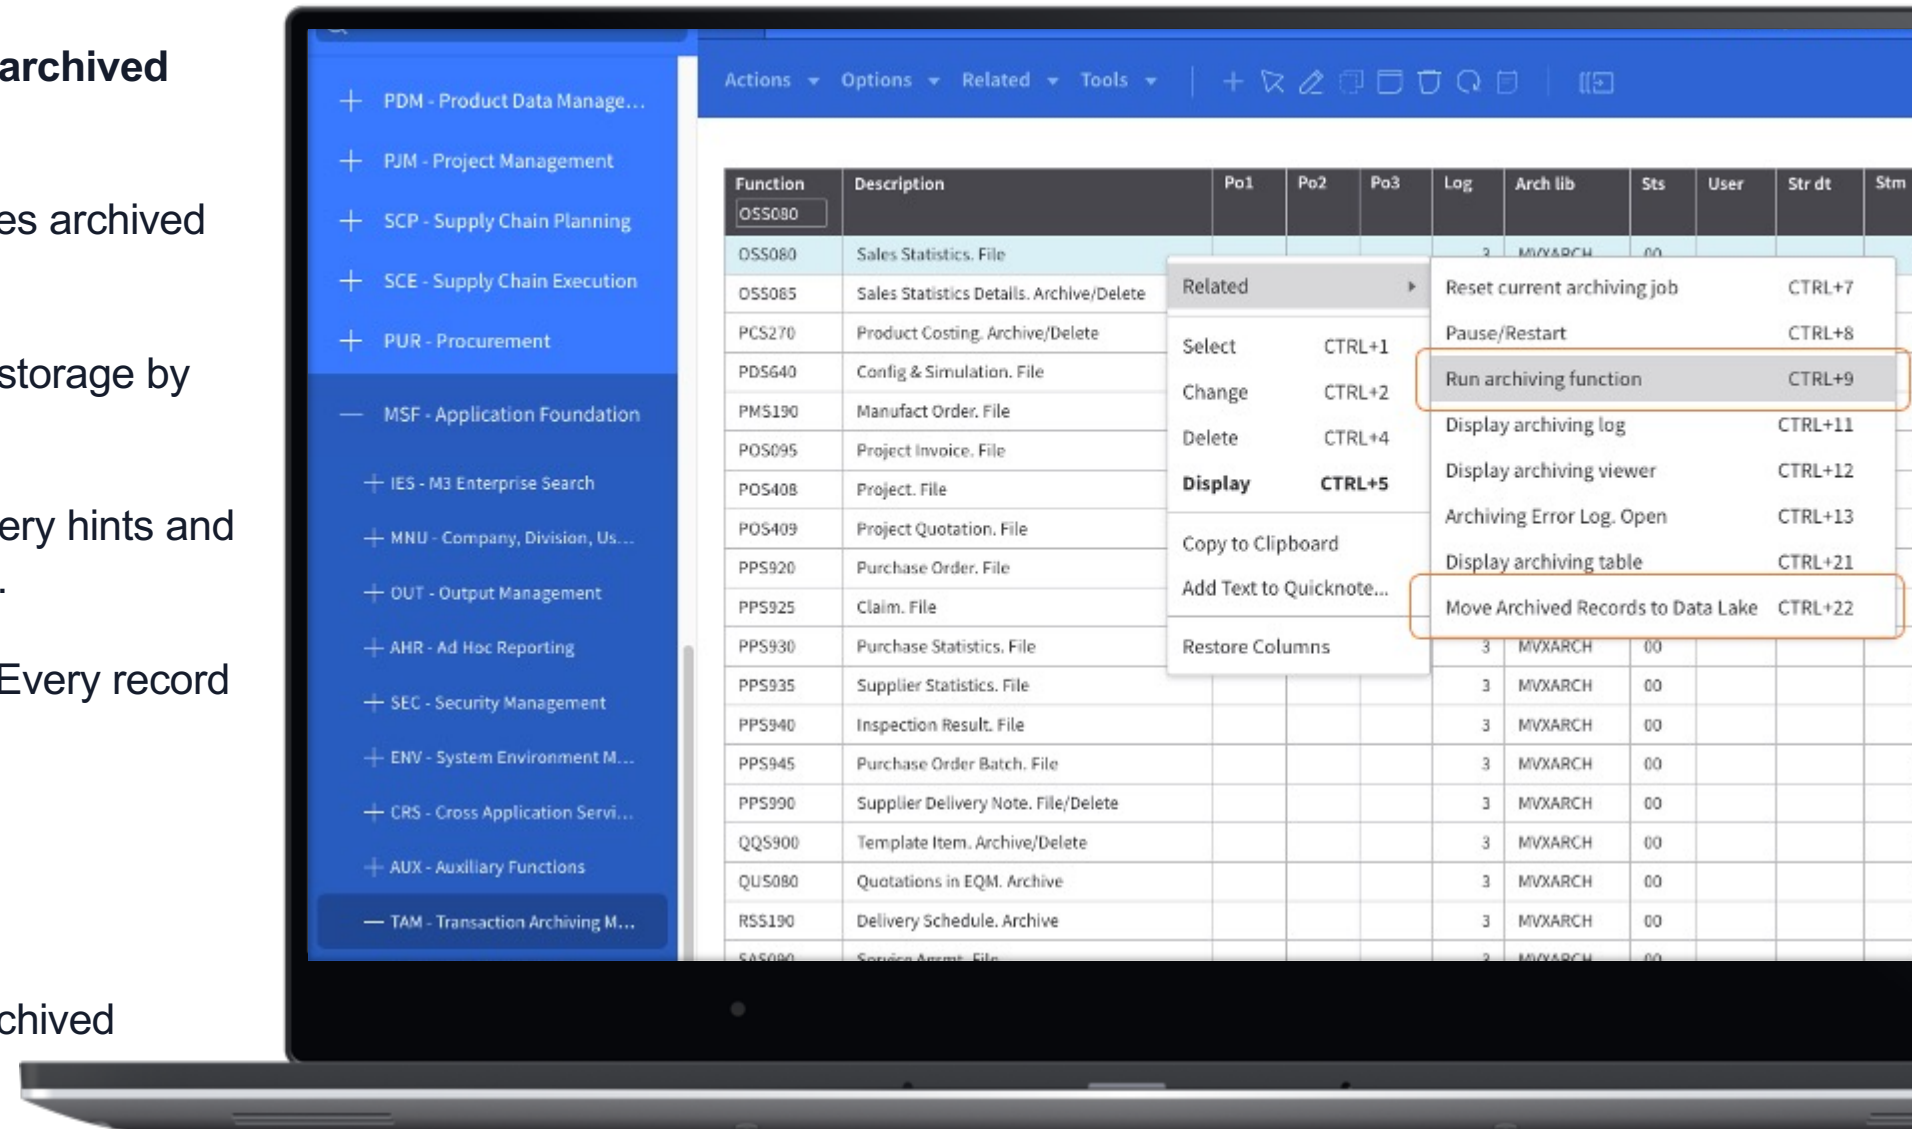
This optional process of archiving is reducing the cost for M3 database storage by moving data into much lower-cost storage.

You can query the archived data in Data Lake by using Infor-specific query hints and present the data in query tools or ETL the data to an on-prem database.

Archived data has a specific flag set that protects it from being purged. Every record in the archived data also has a record-level archive flag set.

Main features are:

- M3 archive functionality works the same
- An additional (optional) step transfers archived data from the M3 archived schema into Data Lake



XtendM3 Dynamic Tables

XtendM3

EXTMAS

Deploy Refresh Disable Export

This table adds the extra info to MOSMAS

Fields Indexes

Fields

Name	Type	Length	Number of d
EXREGN	STRING	18	0
EXCTRY	STRING	32	0
EXCITY	STRING	28	0
EXSTRE	STRING	60	0
EXSNUM	DECIMAL	5	0
EXPERS	STRING	40	0
EXAREA	STRING	250	0

Data Lake Publisher

Subscriptions Configuration

Document Subscriptions

Available	Selected
EXTM	Search Selected
EXTMAS : CUSTOM DATABASE TABLE	ACERTH : Approval Tag 0/(CH)
	ACLHED : Claim Header (AZ)
	CSYFTL : Feature Lines
	CSYFTS : Feature Statuses
	CSYFTT : Feature selections
	CSYFTU : System Update
	OOCHRG : Customer order, charges 1/(OE)
	OODOCU : Customer order, documents 1/(OE)

Additional integrations

- M3 DMP publishing data to Data Fabric using batch ingestion
- Field Audit Trail publishing to Data Lake
- Fashion PLM streaming to Data Lake
- CLM publishing to Data Lake as a source for Analytics.

The screenshot displays the Atlas data management interface. On the left, a search bar is present above a list of data objects. The selected object, **FP_STYLELABEL**, is highlighted. The main panel shows a detailed view of this object, which contains 27 data objects. The table below lists these objects with their Object IDs, Channels, and Indexed dates.

Details	Object ID	Channel	Indexed date
<input type="checkbox"/>	2017-5ab939e3-bf1a-3ae8-8931-328a5916c2f5	Stream	9 okt. 2023 06:30:57
<input type="checkbox"/>	2017-94ee55c1-a12b-3b46-a0a6-cff0ea5238e0	Stream	9 okt. 2023 06:12:26
<input type="checkbox"/>	2017-a6e09e6a-350e-3bf6-9a1d-5e0f15fd9711	Stream	9 okt. 2023 04:50:10
<input type="checkbox"/>	2017-91410b09-d3a1-383e-8ec6-57bc3b151adb	Stream	9 okt. 2023 04:39:20
<input type="checkbox"/>	2017-9b6522db-460a-36ec-a1b0-e5be89be299e	Stream	9 okt. 2023 04:29:10
<input type="checkbox"/>	2017-3f2c6909-2813-35fd-81bd-052eec6b0813	Stream	9 okt. 2023 04:16:14
<input type="checkbox"/>	2017-5ed15595-dbcd-3af4-b560-212e0aeb58cd	Stream	9 okt. 2023 04:03:38
<input type="checkbox"/>	2017-77bc3c82-55b4-3946-b4ba-2baf0957af1f	Stream	9 okt. 2023 03:53:13
<input type="checkbox"/>	2017-47e60172-596a-3d6a-8b71-d1fb8c4006ff	Stream	9 okt. 2023 02:48:39
<input type="checkbox"/>	2017-e15f778d-8629-35ba-b193-4d8aa816a7b6	Stream	9 okt. 2023 02:30:22
<input type="checkbox"/>	2017-1b466fe9-d364-339b-a86f-9713fa201cfb	Stream	9 okt. 2023 02:15:39
<input type="checkbox"/>	2017-1477eaa8-5dce-30ab-ba16-f51450a2597f	Stream	7 okt. 2023 09:45:51
<input type="checkbox"/>	2017-175085b7-b65a-314c-a86a-5073d238da68	Stream	7 okt. 2023 09:28:00
<input type="checkbox"/>	2017-91d77773-5ffe-38a7-be3b-7ea89cb838b6	Stream	7 okt. 2023 00:15:35
<input type="checkbox"/>	2017-03e90f32-fed8-30dd-bb69-790678f0ea94	Stream	6 okt. 2023 19:23:45
<input type="checkbox"/>	2017-b77e2e0e-0f63-31f6-9e72-0a4a6df78eb9	Stream	6 okt. 2023 19:13:14
<input type="checkbox"/>	2017-d04c5504-949c-3e37-8f09-13f8017cad34	Stream	6 okt. 2023 18:58:28
<input type="checkbox"/>	2017-850510f0-40fb-37b0-b71b-6fd5569f6177	Stream	6 okt. 2023 18:47:50
<input type="checkbox"/>	2017-73722ab8-8041-3d97-8b00-9aebb261f6b6	Stream	6 okt. 2023 18:32:42
<input type="checkbox"/>	2017-15378e29-3004-3f8c-a6d9-77aa5e6e8b6f	Stream	6 okt. 2023 18:21:59
<input type="checkbox"/>	2017-c550b4c4-7d1c-3c5d-9c0b-786c80822f49	Stream	6 okt. 2023 12:44:28
<input type="checkbox"/>	2017-cda1bd24-acde-371d-97d5-0f14d95a7a4a	Stream	2 okt. 2023 16:50:10
<input type="checkbox"/>	2017-f15536c5-fc6f-33e4-99d9-6e8c9a620fef	Stream	2 okt. 2023 16:40:14

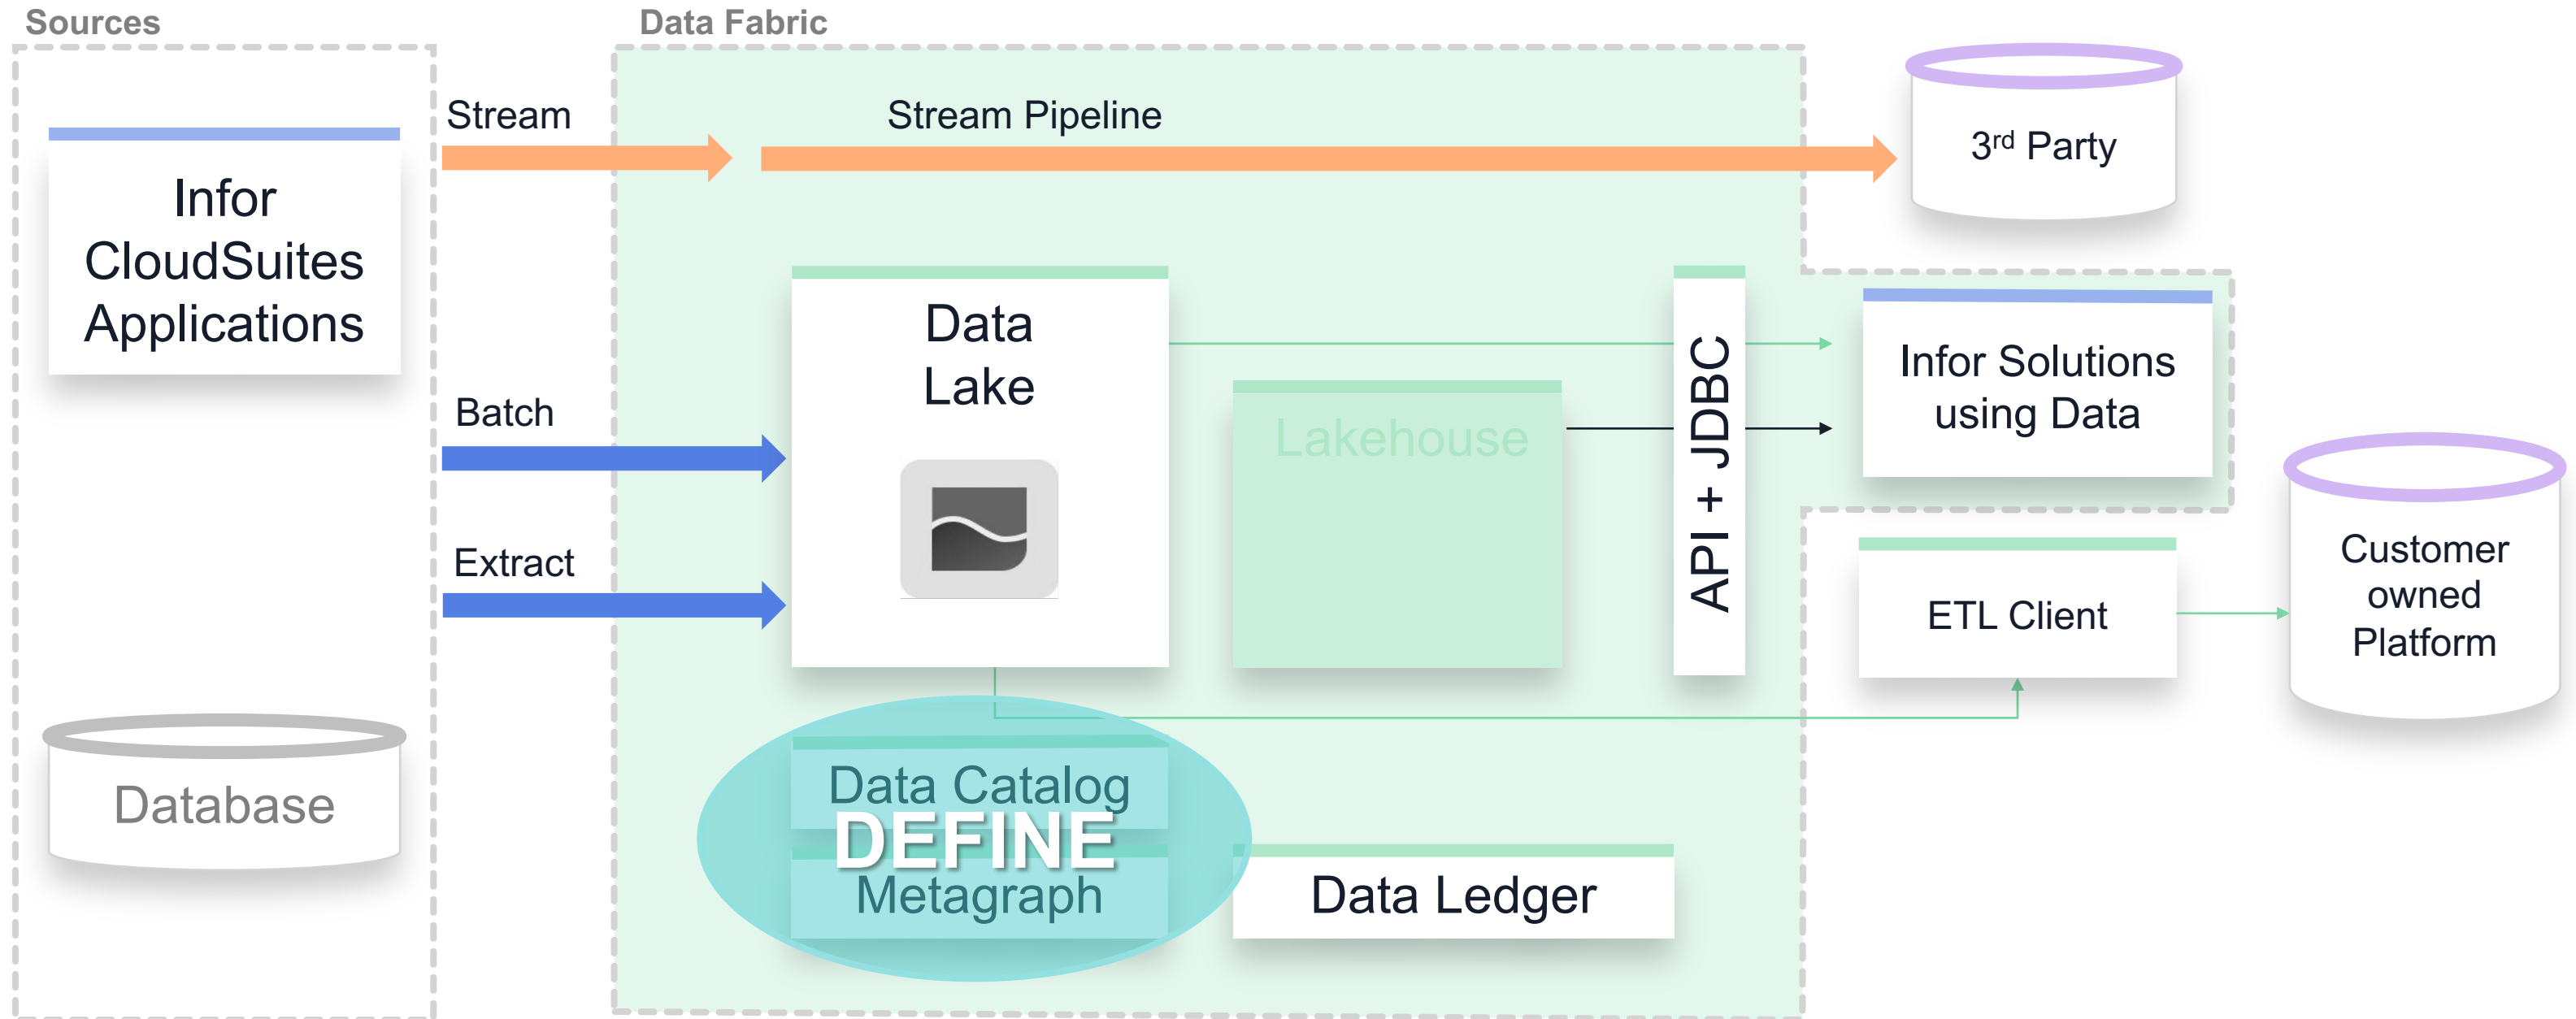
On the right side of the interface, there are several action buttons: **Purge** (red), **Mark corrupt** (blue), and **Download** (blue). The interface also shows a search bar, a list of data objects, and a detailed view of the selected object.



04

Metadata – Defining what's in Data Fabric

Data Fabric – Define Data



Data Catalog

The screenshot displays the M3 Data Catalog interface. On the left is a navigation sidebar with options like Home, Connect, Monitors & Workflows, Scripting, OneView, Data Catalog, Object Schemas, Data Lake Views, Lakehouse Schemas, Schema Extensions, Locale Selections, Configuration, Configuration Management, Authorizations, and About ION. The main area is titled 'Object Schemas' and shows 528 objects. A search filter 'Filter by name' is present. Below the filter are several object cards for JSON schemas: EXTWHS, FAPIBH, MITAFS, AnyDocument, and MITDCU. The 'Object Schema CFACIL' is selected, showing its details: Name (CFACIL), Title (Facility master), Description (Facility master), Type (JSON), Subtype (Newline-delimited), and Library (Custom). A 'Formatted View' tab is active, displaying a table of properties.

Property	Indicators	Title	Description	Data Type	Required	Position	Length	Digits Before Decimal	Digits After Decimal
CONO		company	company	Number	<input checked="" type="checkbox"/>	1		3	0
FACI		facility	facility	String	<input checked="" type="checkbox"/>	2	3		
FACN		name	name	String	<input type="checkbox"/>	3	30		
DIVI		division	division	String	<input type="checkbox"/>	4	3		
WHLO		main warehouse	main warehouse	String	<input type="checkbox"/>	5	3		
ACGR		object access group	object access group	String	<input type="checkbox"/>	6	10		
PYAD		our invoicing address	our invoicing address	String	<input type="checkbox"/>	7	3		
CUNO		customer	customer	String	<input type="checkbox"/>	8	10		
SUNO		supplier	supplier	String	<input type="checkbox"/>	9	10		
SHPL		shift planning	shift planning	Number	<input type="checkbox"/>	10		1	0
GFAC		global facility	global facility	String	<input type="checkbox"/>	11	3		
APSA		aps scheduling enabled	aps scheduling enabled	Number	<input type="checkbox"/>	12		1	0
CCMT		shift pattern adjustments	shift pattern adjustments	Number	<input type="checkbox"/>	13		1	0
TSDA		swb timestamp	swb timestamp	Number	<input type="checkbox"/>	14		8	0
TSTE		swb timestamp	swb timestamp	Number	<input type="checkbox"/>	15		6	0
APSS		swb in progress	swb in progress	Number	<input type="checkbox"/>	16		1	0
SCHH		scheduling horizon	scheduling horizon	Number	<input type="checkbox"/>	17		3	0
TXID		text identity	text identity	Number	<input type="checkbox"/>	18		13	0
RGDT		entry date	entry date	Number	<input type="checkbox"/>	19		8	0
RGTM		entry time	entry time	Number	<input type="checkbox"/>	20		6	0
LMDT		change date	change date	Number	<input type="checkbox"/>	21		8	0
CHNO		change number	change number	Number	<input type="checkbox"/>	22		3	0

Metagraph

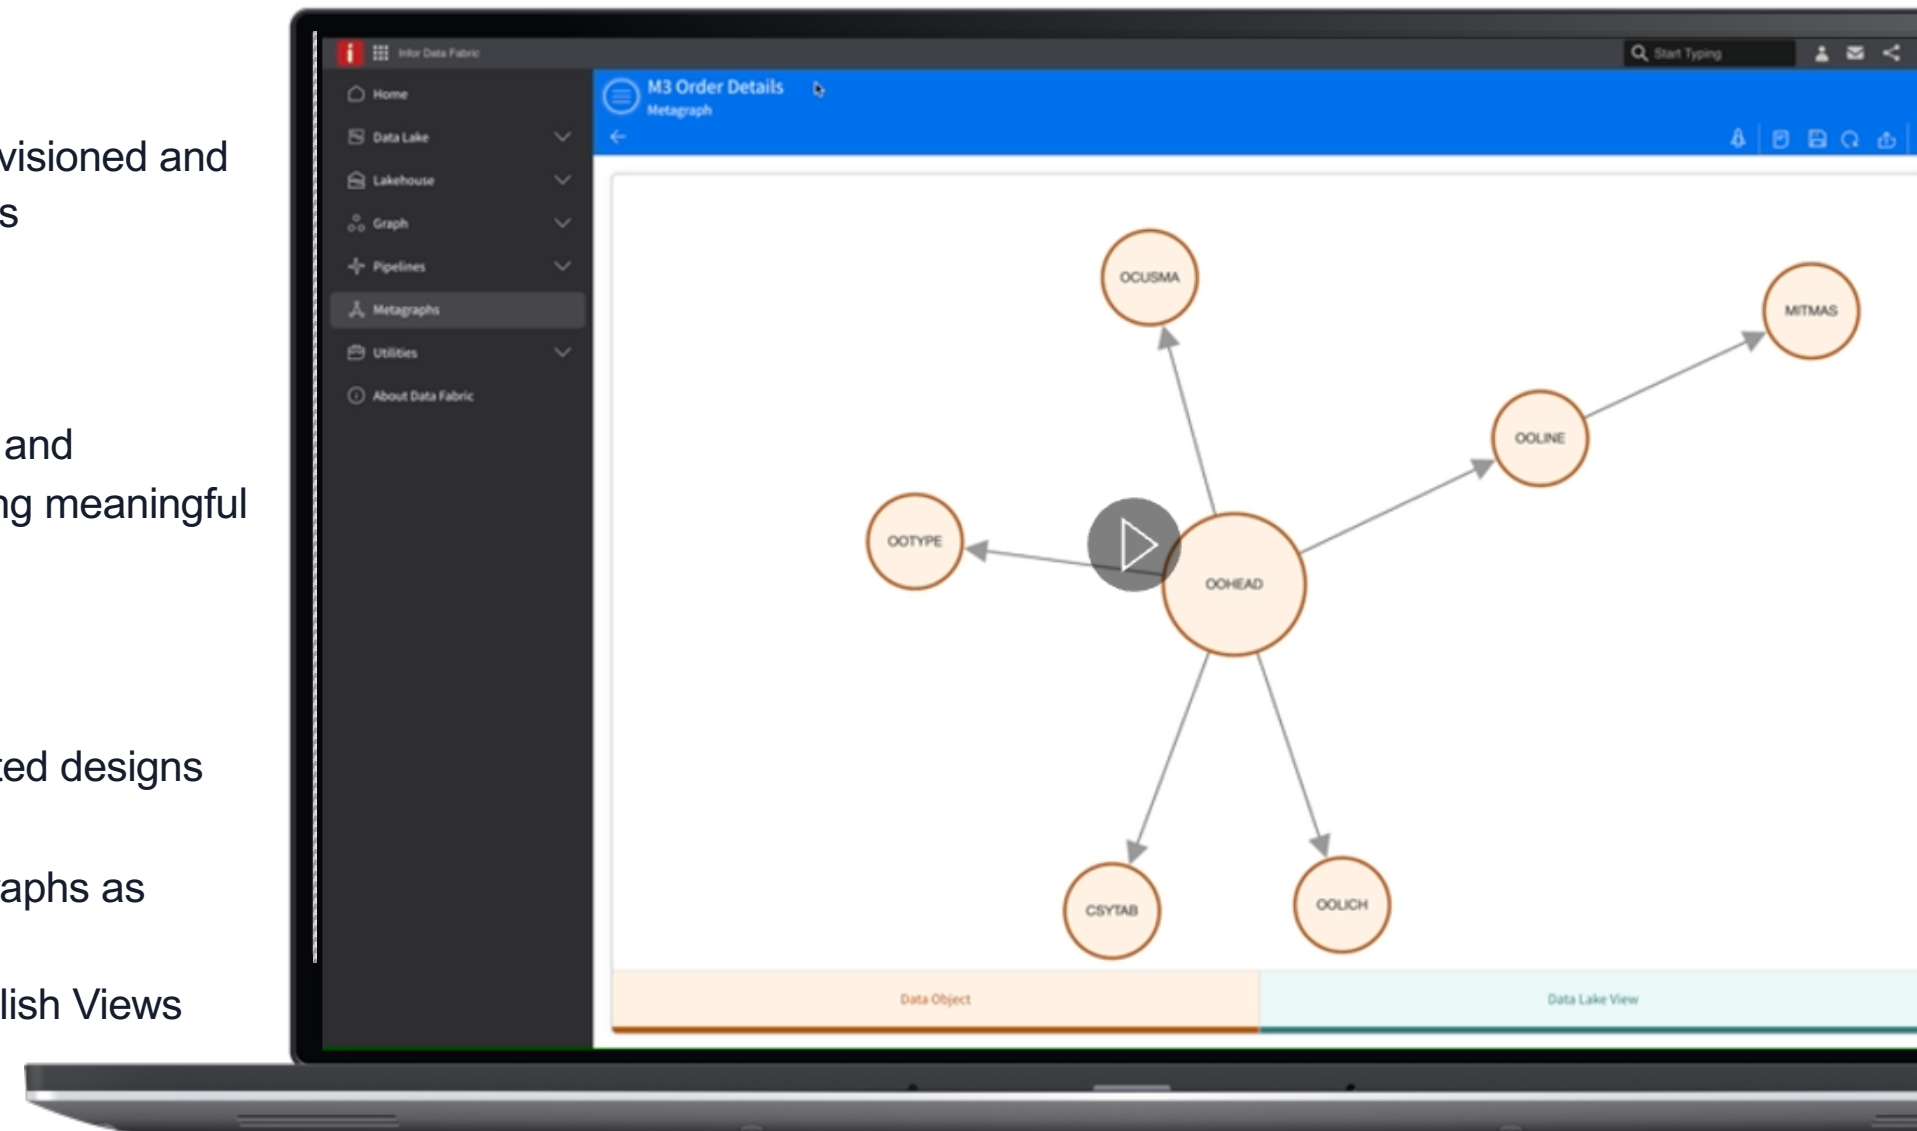
Accelerating the content development curve.

Metagraphs alleviate headaches and reduces guesswork with Infor-provisioned and user-created Metagraphs that provide domain-specific canonical models representing functional application modules, screens, and reports.

Graph-led representation of your metadata footprint is easily converted and published as a Data Lake View to accelerate developer time in producing meaningful content.

Main features are:

- Better understand your data at the metadata level with graph-oriented designs describing relationships
- Drag & drop modeling allows users to compose and extend Metagraphs as domains evolve
- Operationalize your Metagraph with the Publishing Wizard and publish Views straight to Data Lake within 3 steps



What is Metagraphs?

Modelling tool
Metagraphs can be used to model relations between data objects in Data Lake



Data fabric
Tool within Data Fabric, using data published to Data Lake



SQL-views
When building models – SQL views are automatically generated

M3BE overview
Models can be used to visualize table connections in different M3BE processes



Analytics
SQL queries can be used to extract data, when creating analytical reports

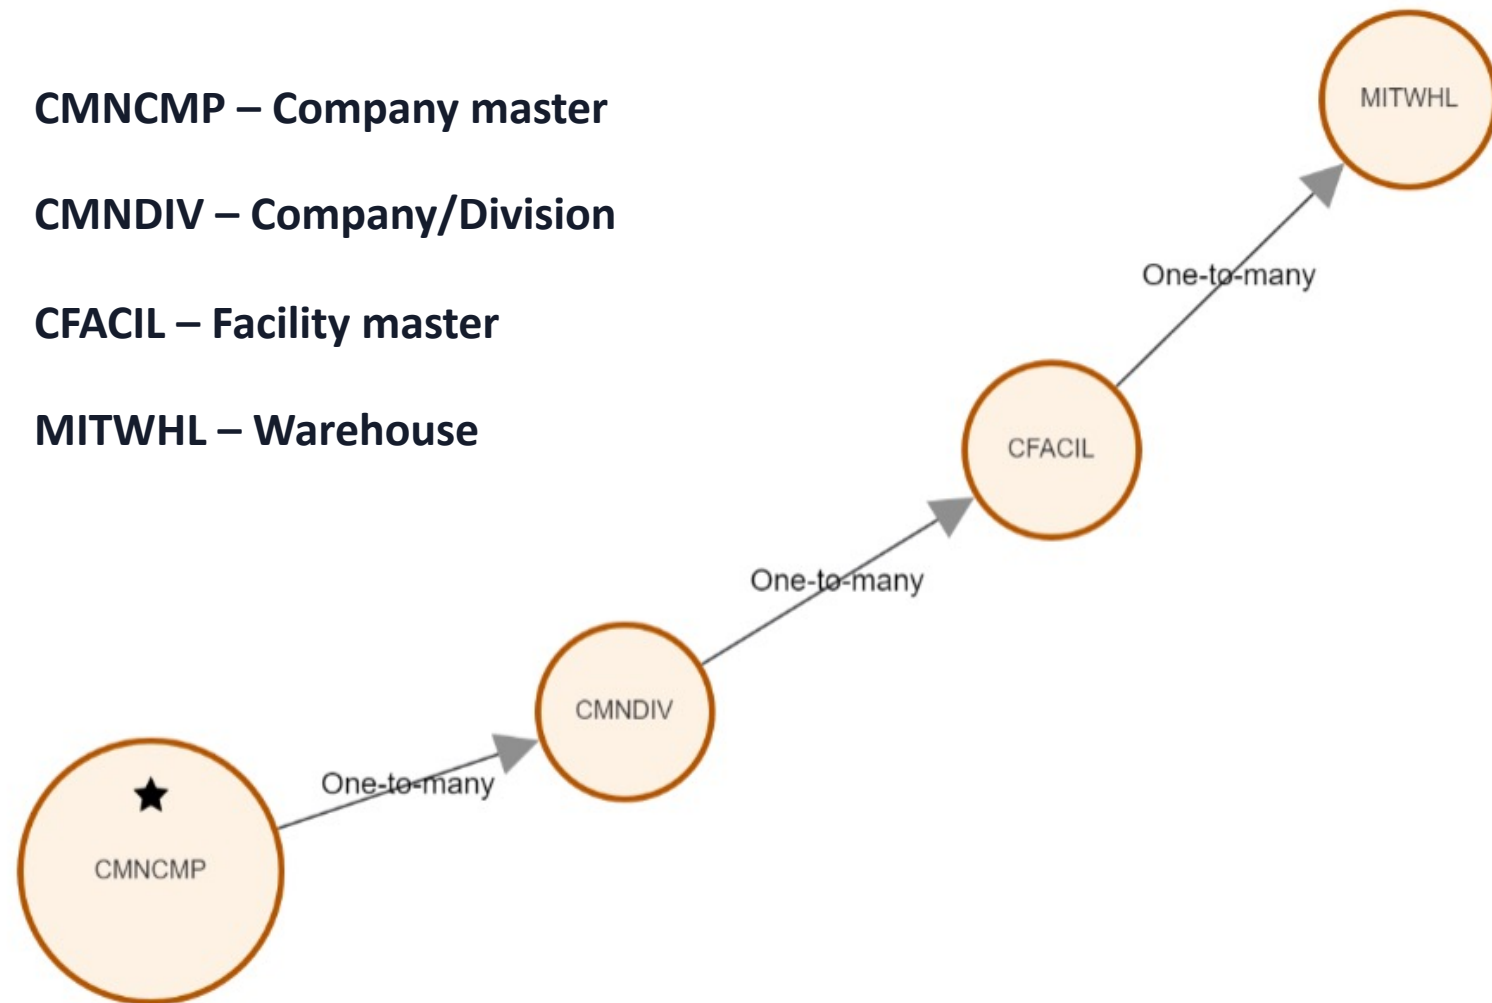
Example

CMNCMP – Company master

CMNDIV – Company/Division

CFACIL – Facility master

MITWHL – Warehouse

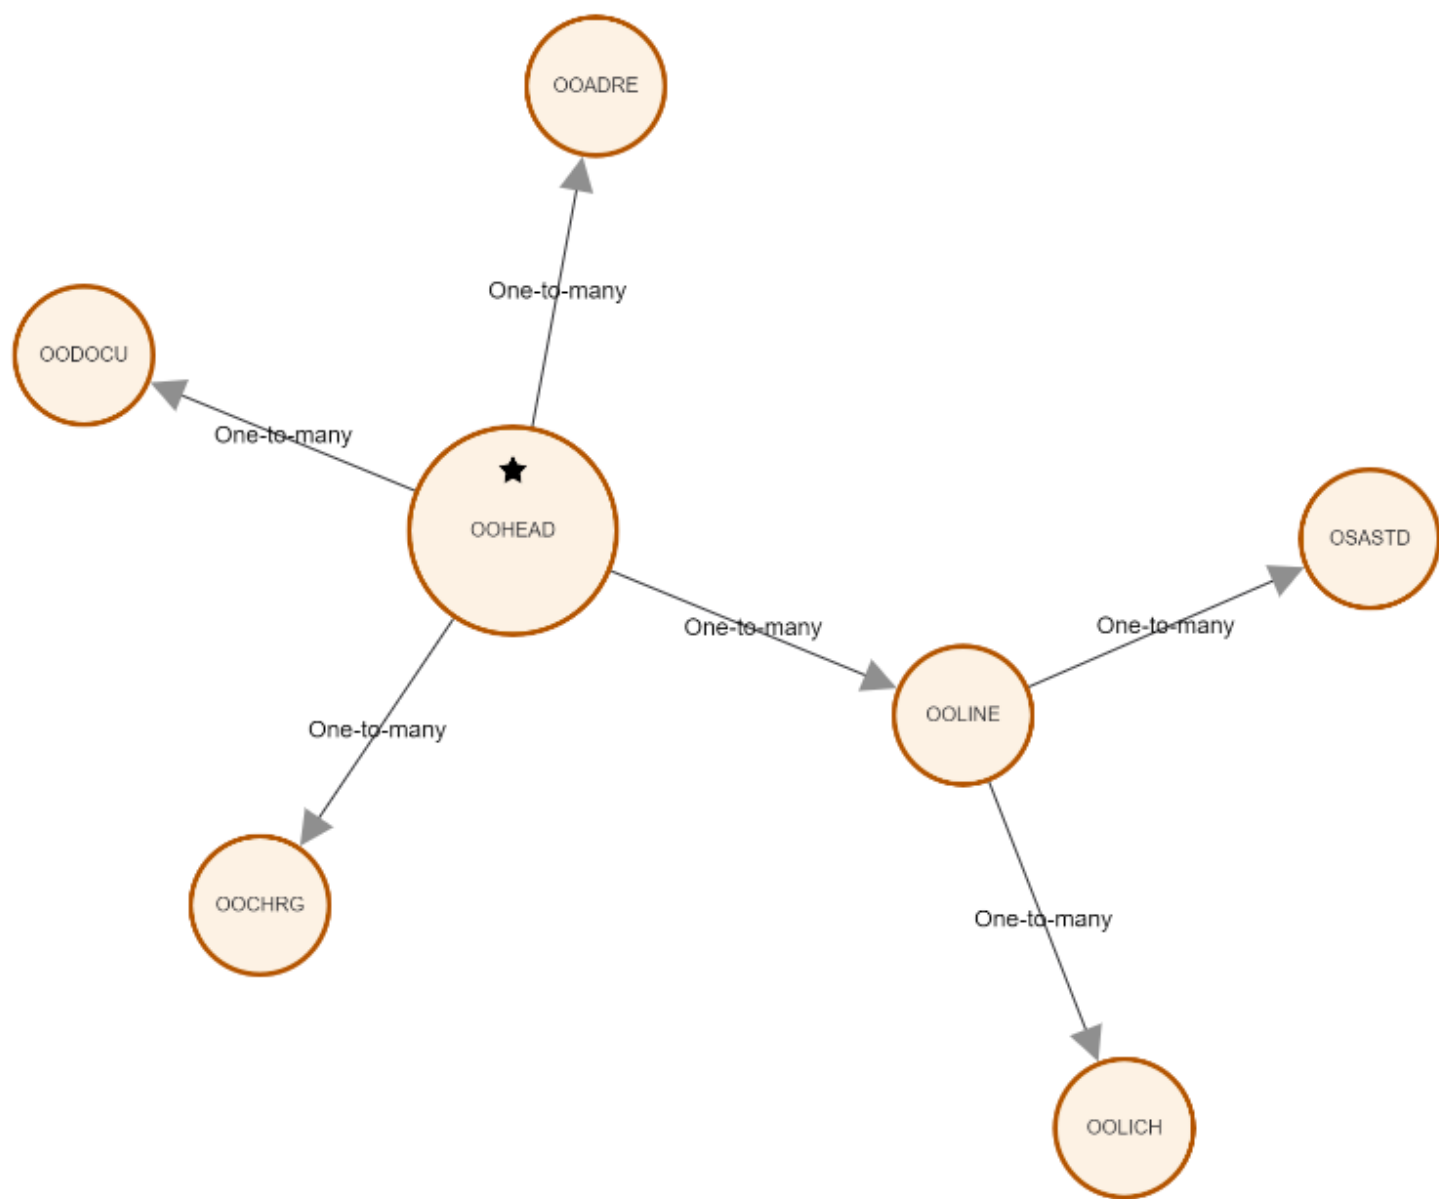


Publish view to Data Lake

View name* Property selection* 3 Generate SQL Statement

Use titles as aliases for view Preview Copy

```
1 SELECT
2   "MITWHL"."WHLO",
3   "CMNCMP"."CONO",
4   "CMNDIV"."DIVI",
5   "CFACIL"."FACI"
6 FROM
7   "CMNCMP"
8   INNER JOIN "CMNDIV" ON "CMNCMP"."CONO" = "CMNDIV"."CONO"
9   INNER JOIN "CFACIL" ON "CMNDIV"."CONO" = "CFACIL"."CONO"
10  AND "CMNDIV"."DIVI" = "CFACIL"."DIVI"
11  INNER JOIN "MITWHL" ON "CFACIL"."DIVI" = "MITWHL"."DIVI"
12  AND "CFACIL"."FACI" = "MITWHL"."FACI"
13  AND "CFACIL"."CONO" = "MITWHL"."CONO"
```



```

1  SELECT
2  "OOADRE"."ADID",
3  "OOADRE"."ADRT",
4  "ODOCU"."DOVA",
5  "ODOCU"."DONR",
6  "OOHEAD"."CONO",
7  "OOHEAD"."ORNO",
8  "OOHEAD"."CUNO",
9  "OOLINE"."PONR",
10 "OOLINE"."ITNO",
11 "OOLINE"."ITDS",
12 "OOLICH"."CRID",
13 "OSASTD"."OSSQ",
14 "OOCHRG"."CRID"
15 FROM
16 "OOHEAD"
17 LEFT OUTER JOIN "ODOCU" ON "OOHEAD"."ORNO" = "ODOCU"."ORNO"
18 AND "OOHEAD"."CONO" = "ODOCU"."CONO"
19 LEFT OUTER JOIN "OOCHRG" ON "OOHEAD"."CONO" = "OOCHRG"."CONO"
20 AND "OOHEAD"."ORNO" = "OOCHRG"."ORNO"
21 LEFT OUTER JOIN "OOLINE" ON "OOHEAD"."CONO" = "OOLINE"."CONO"
22 AND "OOHEAD"."ORNO" = "OOLINE"."ORNO"
23 LEFT OUTER JOIN "OOADRE" ON "OOHEAD"."ORNO" = "OOADRE"."ORNO"
24 AND "OOHEAD"."CONO" = "OOADRE"."CONO"
25 LEFT OUTER JOIN "OOLICH" ON "OOLINE"."POSX" = "OOLICH"."POSX"
26 AND "OOLINE"."ORNO" = "OOLICH"."ORNO"
27 AND "OOLINE"."PONR" = "OOLICH"."PONR"
28 AND "OOLINE"."CONO" = "OOLICH"."CONO"
29 LEFT OUTER JOIN "OSASTD" ON "OOLINE"."POSX" = "OSASTD"."POSX"
30 AND "OOLINE"."CONO" = "OSASTD"."CONO"
31 AND "OOLINE"."ORNO" = "OSASTD"."ORNO"
32 AND "OOLINE"."PONR" = "OSASTD"."PONR"
  
```



```

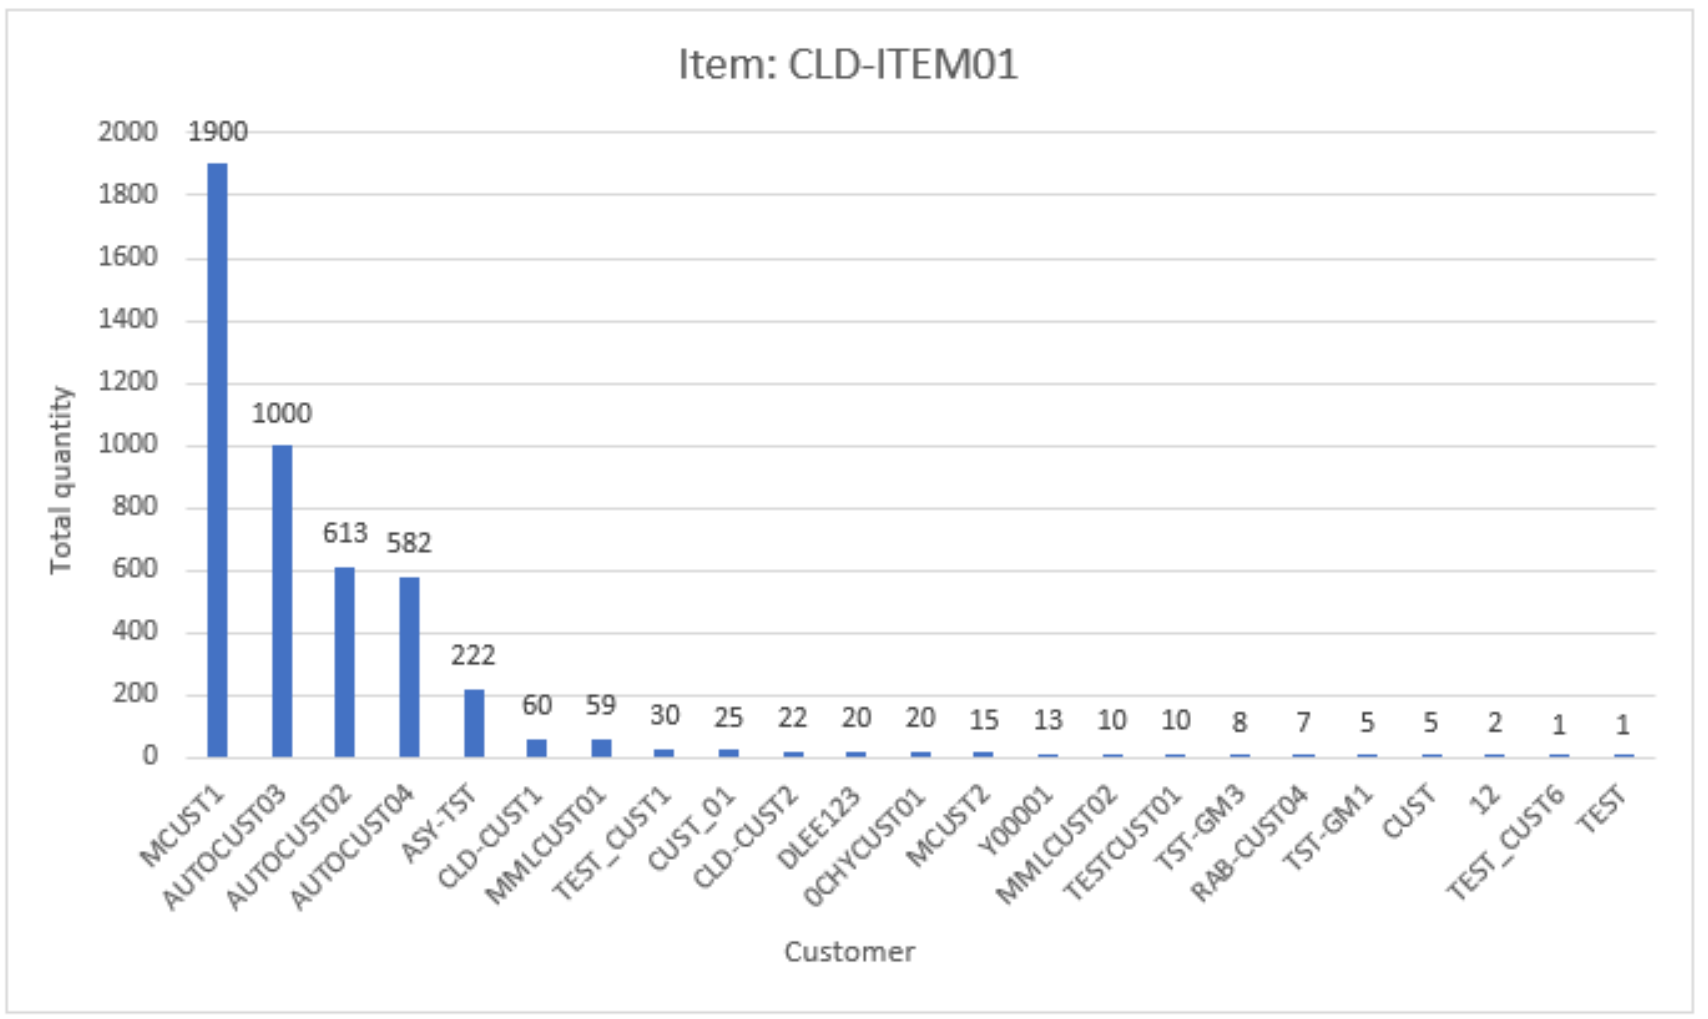
SELECT
  "OOHEAD"."CONO" AS "COMPANY",
  "OOHEAD"."CUNO" AS "CUSTOMER NO",
  "OOLINE"."ITNO" AS "ITEM NO",
  SUM("OOLINE"."ORQT") AS "QUANTITY"
FROM
  "OOHEAD"
  LEFT OUTER JOIN "OOLINE" ON "OOHEAD"."CONO" = "OOLINE"."CONO"
  AND "OOHEAD"."ORNO" = "OOLINE"."ORNO"
WHERE
  "OOHEAD"."CONO" = '770' AND "OOLINE"."ITNO" = 'CLD-ITEM01'
GROUP BY
  "OOHEAD"."CONO", "OOHEAD"."CUNO", "OOLINE"."ITNO"
ORDER BY
  "QUANTITY" DESC

```

Results 1 Results 2

SELECT "OOHEAD"."CONO" AS "COMPANY", "OOHEAD"."CUNO" AS "CUSTOMER NO", "OOLINE"."ITNO" AS "ITEM" Enter a St

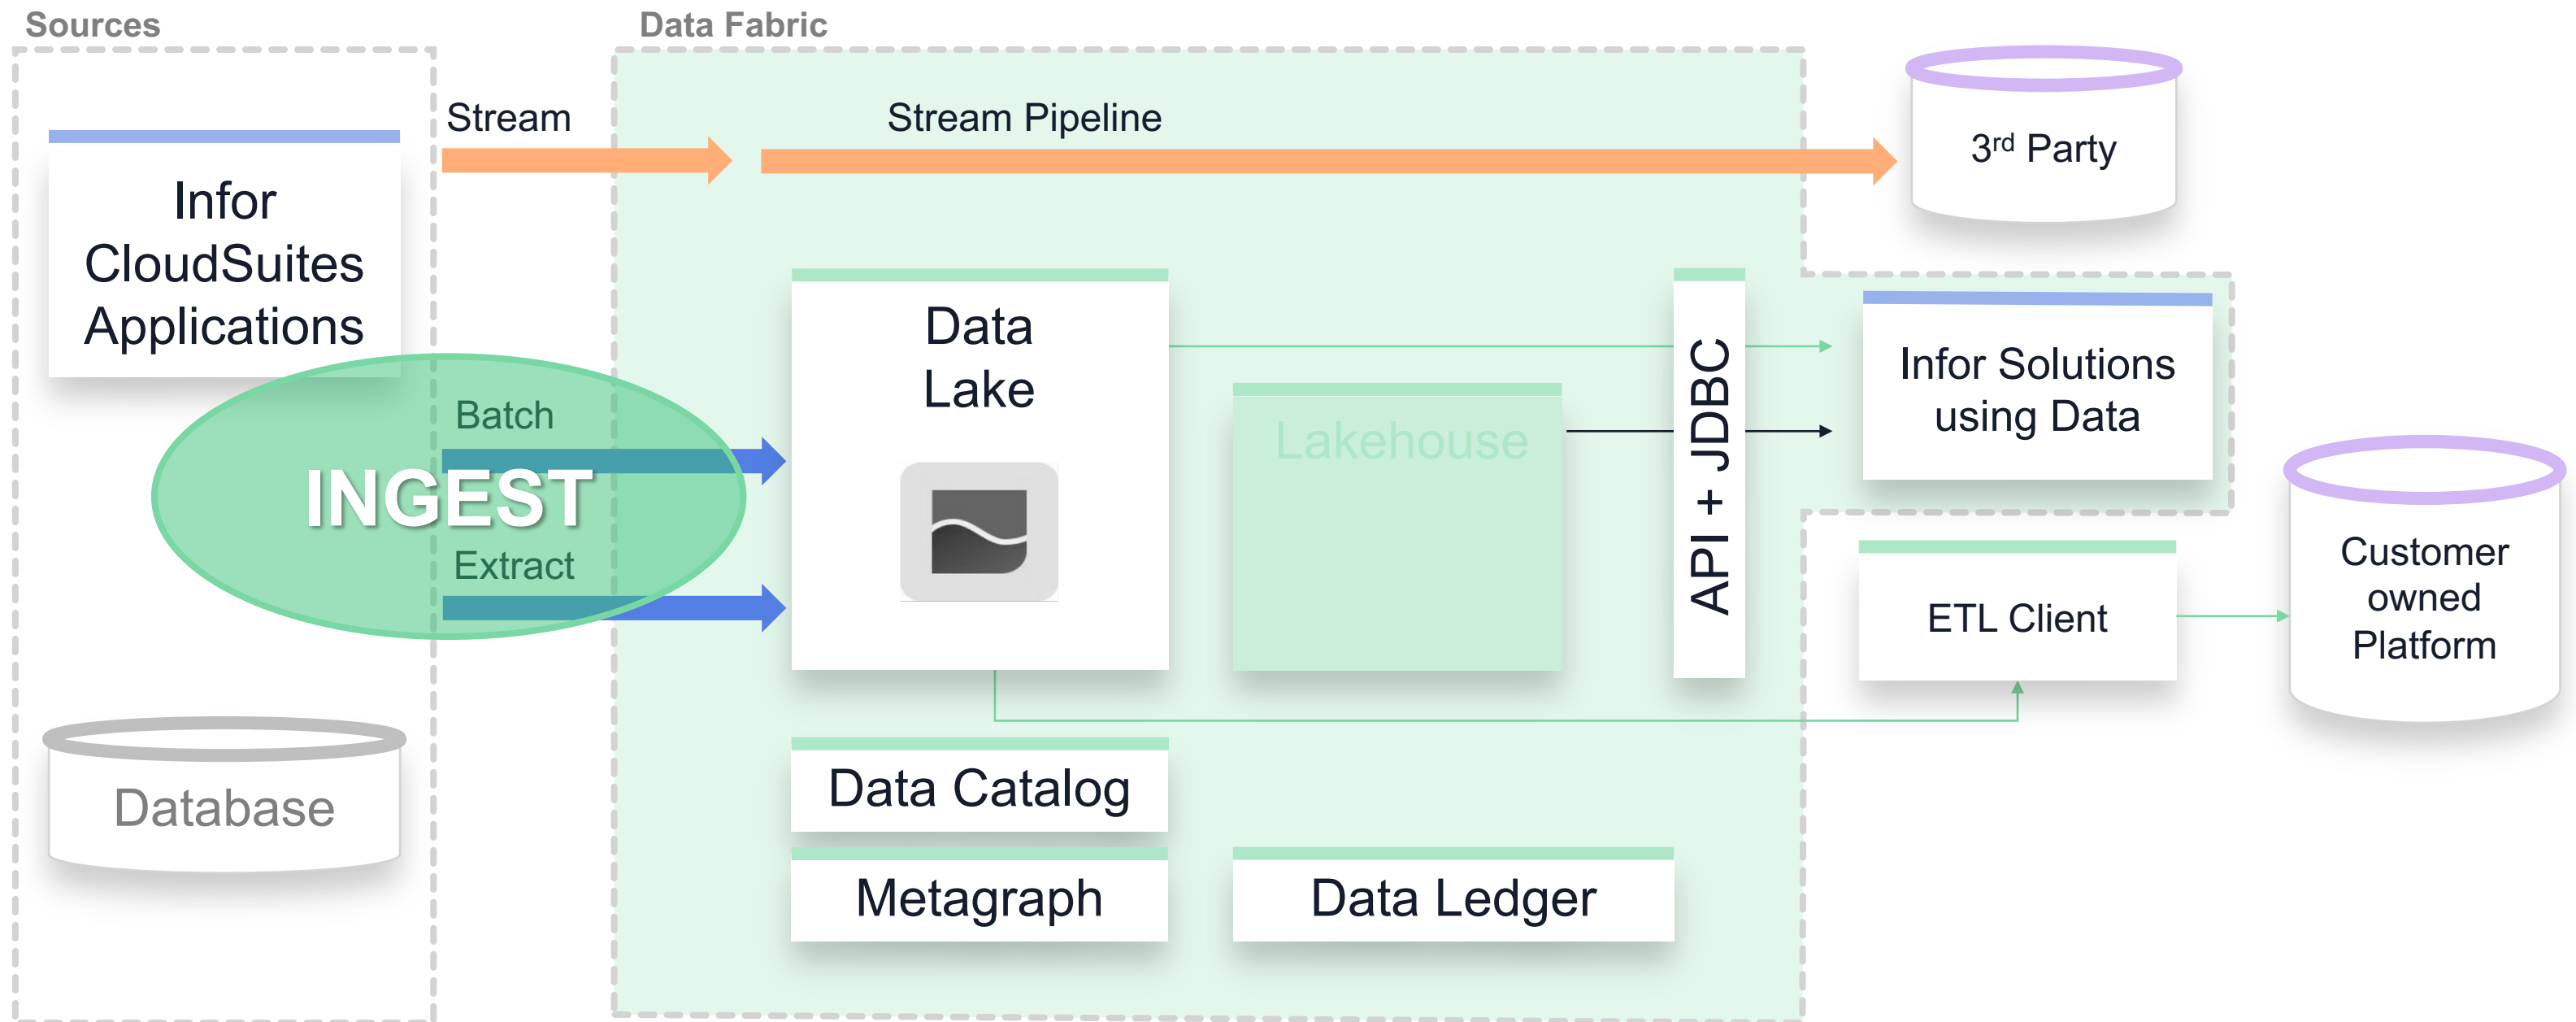
	123 COMPANY	ABC CUSTOMER NO	ABC ITEM NO	123 QUANTITY
1	770	MCUST1	CLD-ITEM01	1900
2	770	AUTOCUST03	CLD-ITEM01	1000
3	770	AUTOCUST02	CLD-ITEM01	613
4	770	AUTOCUST04	CLD-ITEM01	582
5	770	ASY-TST	CLD-ITEM01	222
6	770	CLD-CUST1	CLD-ITEM01	60
7	770	MMLCUST01	CLD-ITEM01	59
8	770	TEST_CUST1	CLD-ITEM01	30
9	770	CUST_01	CLD-ITEM01	25
10	770	CLD-CUST2	CLD-ITEM01	22
11	770	DLEE123	CLD-ITEM01	20
12	770	OCHYCUST01	CLD-ITEM01	20
13	770	MCUST2	CLD-ITEM01	15
14	770	Y00001	CLD-ITEM01	13
15	770	MMLCUST02	CLD-ITEM01	10
16	770	TESTCUST01	CLD-ITEM01	10
17	770	TST-GM3	CLD-ITEM01	8
18	770	RAB-CUST04	CLD-ITEM01	7
19	770	CUST	CLD-ITEM01	5
20	770	TST-GM1	CLD-ITEM01	5
21	770	12	CLD-ITEM01	2
22	770	TEST_CUST6	CLD-ITEM01	1
23	770	TEST	CLD-ITEM01	1





05 Ingestion – Getting data into Data Fabric

Data Fabric – Ingest Data



M3 Data Lake Publisher

The screenshot displays the M3 Data Lake Publisher interface. At the top, there is a navigation bar with the Infor M3 logo, a search bar containing 'Start Typing', and user information for 'Per Melander'. Below this is a secondary navigation bar with 'Data Lake Publisher' and a hamburger menu icon. A left-hand sidebar lists various application areas, with 'Data Lake Publisher' currently selected. The main content area is titled 'Data Lake Publisher' and contains two tabs: 'Subscriptions' (active) and 'Configuration'. The 'Subscriptions' tab shows a 'Document Subscriptions' window with two columns: 'Available' and 'Selected'. Each column has a search bar and a list of document subscriptions. The 'Available' list includes items like 'ACLHED : Claim Header (AZ)', 'ACLSPC : Claim Spec (AZ)', and 'ACREWI : Crew Information 0/(FW)'. The 'Selected' list includes 'ACERTH : Approval Tag 0/(CH)', 'CSYCAL : System calendar', and 'OCUSMA : Customer', with 'OCUSMA : Customer' highlighted. Navigation icons for save, refresh, and back are visible in the top right of the subscription window.

Available	Selected
ACLHED : Claim Header (AZ)	ACERTH : Approval Tag 0/(CH)
ACLSPC : Claim Spec (AZ)	CSYCAL : System calendar
ACLTYP : Claim Type	MITMAS : Item Master
ACLVER : Weight and Volume Information	OCUSMA : Customer
ACREWI : Crew Information 0/(FW)	OOHEAD : CO header file
ACSALS : Simulation Agreement summary 0/(SS)	OOLINE : CO line file
ACUACC : Invoice accounting 1/(UT)	
ACUACL : Agreement line charges	

M3 and Data Fabric

M3 Streaming Toggle

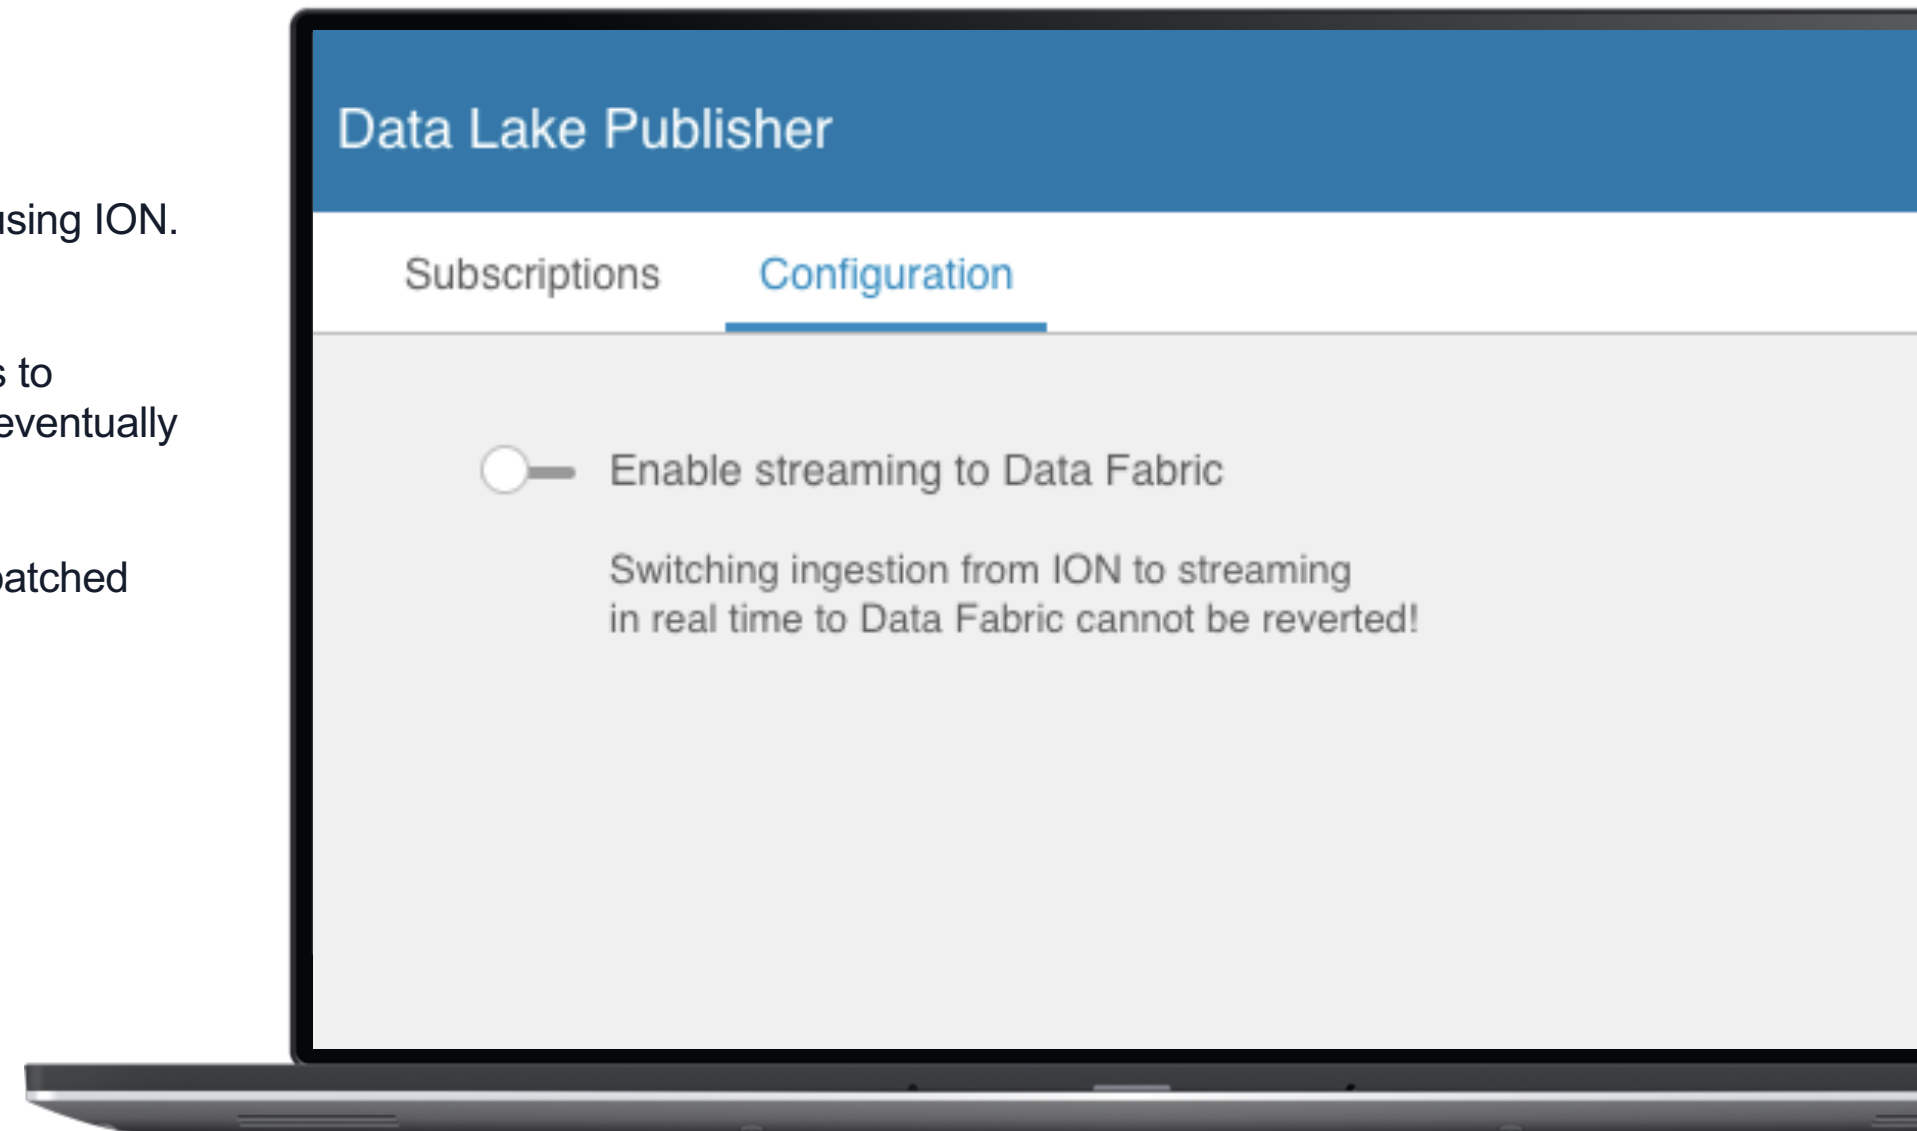
Streaming is a new way to publish data to Data Fabric directly without using ION. When switching to Streaming, you cannot go back to ION!

M3 Streaming is intended to be used with Data Fabric Stream Pipelines to support operational data solutions with real-time requirements. Data is eventually micro-batched by Data Fabric into Data Lake after 10 minutes or 5MB.

Without Stream Pipelines, data will be streamed and eventually micro-batched the same way.

- + **Support real-time solutions with Stream Pipelines**
- + **Data Lake files get an optimized structure**

- **10 min latency in Data Lake**

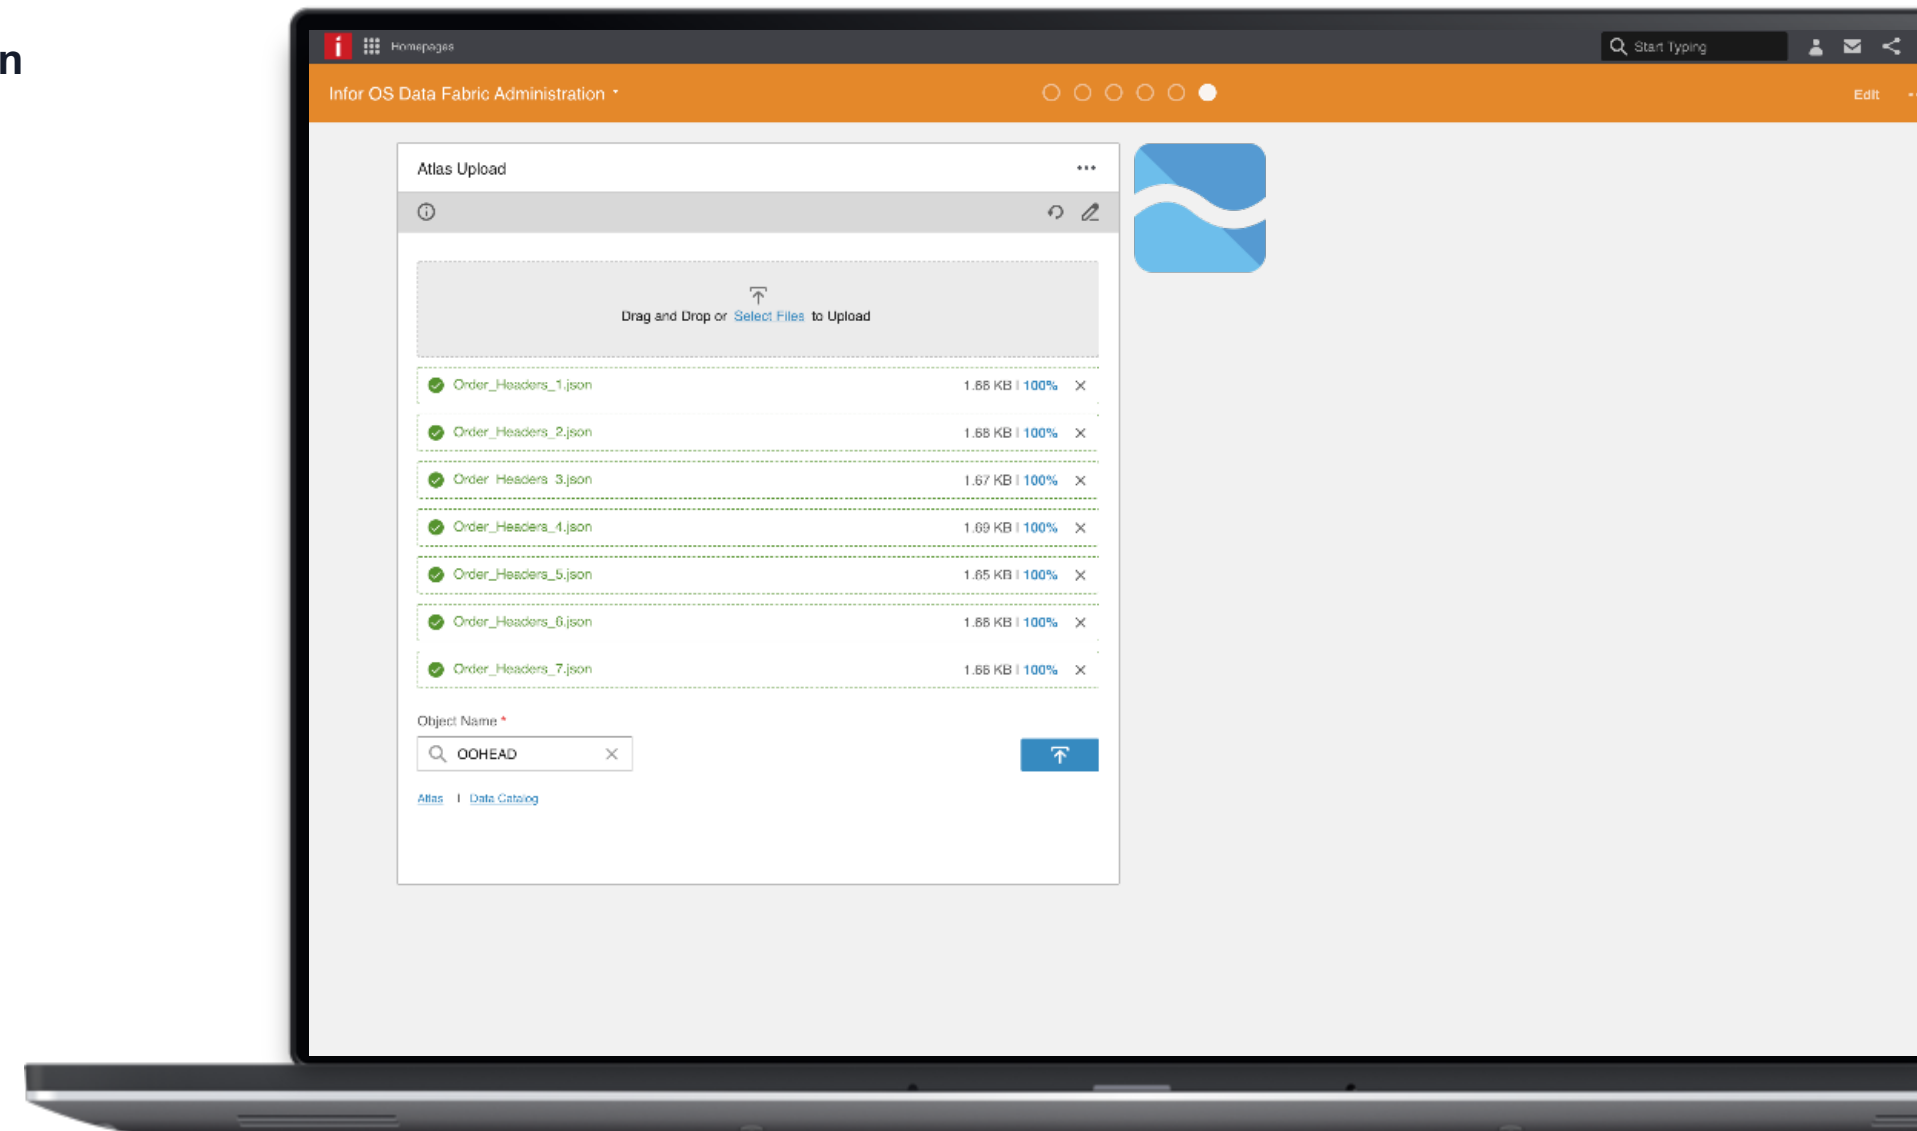


Atlas Upload Widget

Quick & easy method for users to upload data in Data Lake using an Infor OS widget interface.

Transfer local files to Data Lake without setting up data flows to extract from a source or having to integrate with the Data Fabric Ingestion APIs.

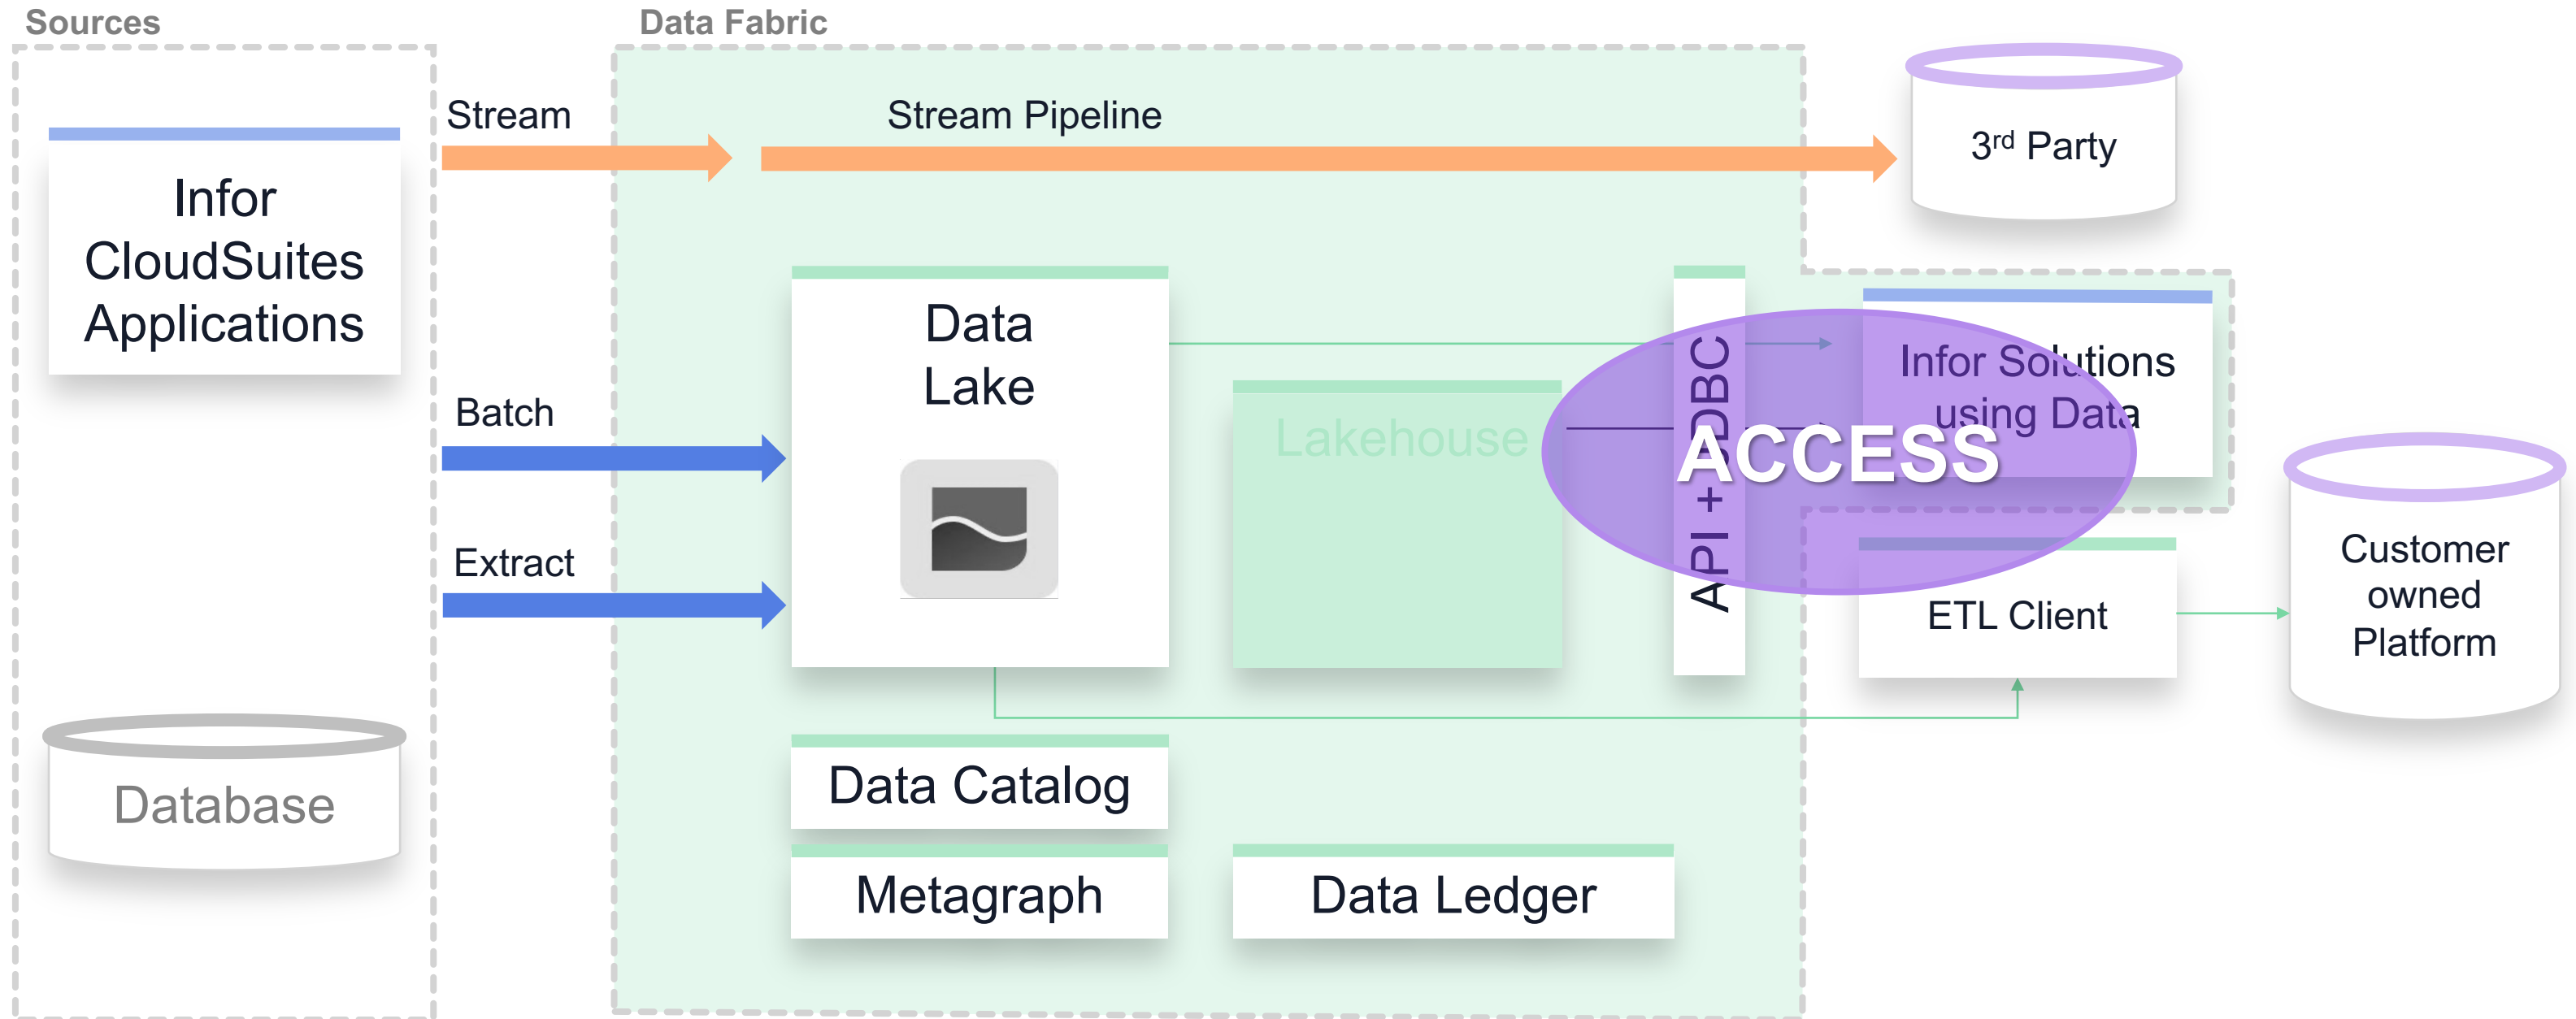
- Available in Ming.le Homepages and Infor Portal (V2)
- Empowers non-technical personas to contribute to Data Lake
- Deploy data from unintegrated data sources



06 How to access data in Data Lake



Data Fabric – Access Data



Compass

Access Data Lake using SQL with variation logic

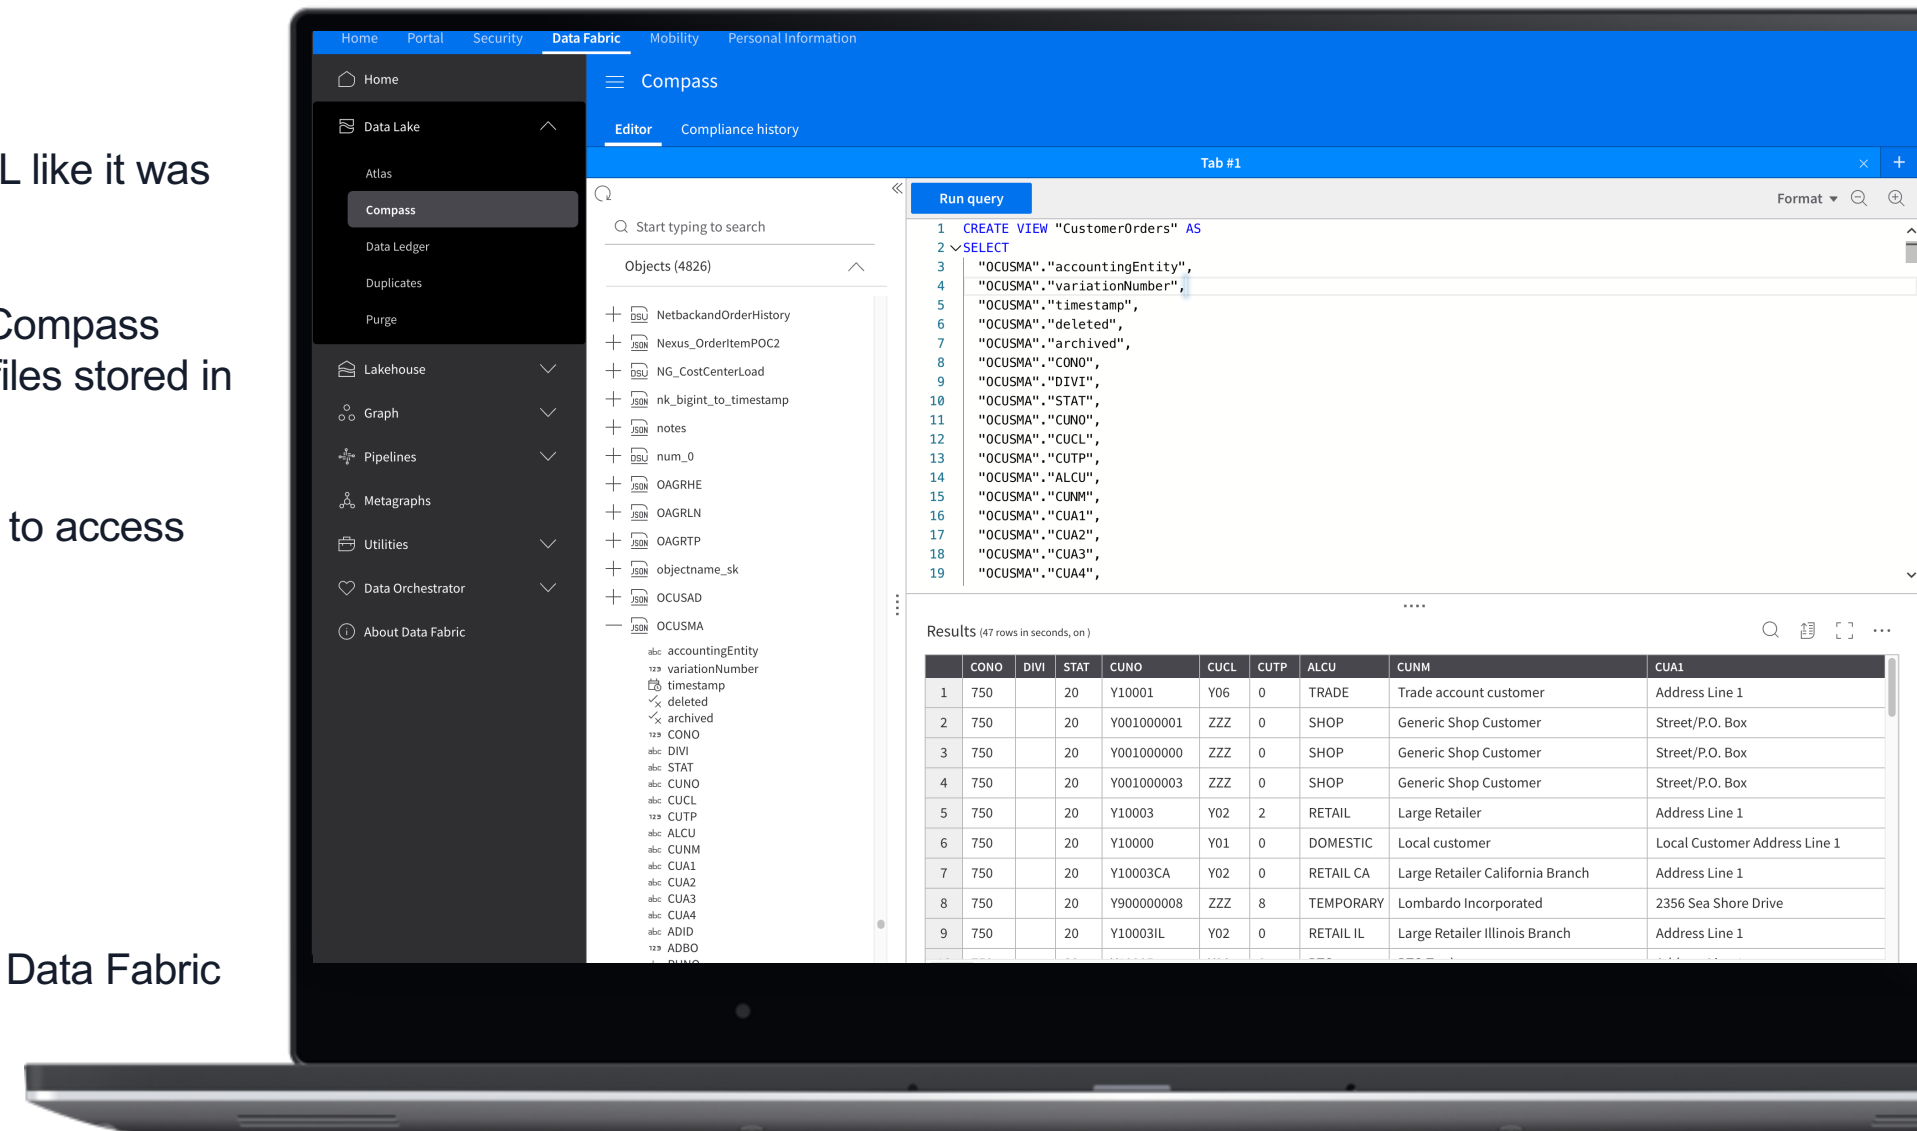
Compass on top of Data Lake gives you full control over data using SQL like it was your transactional M3 database.

Data Lake is not and has not have the performance of a database but Compass caches data in a structure to support high-performance queries on the files stored in Data Lake.

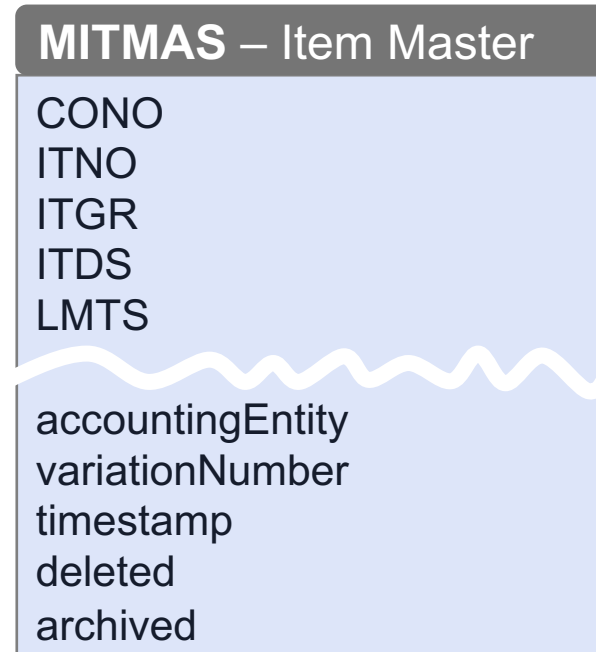
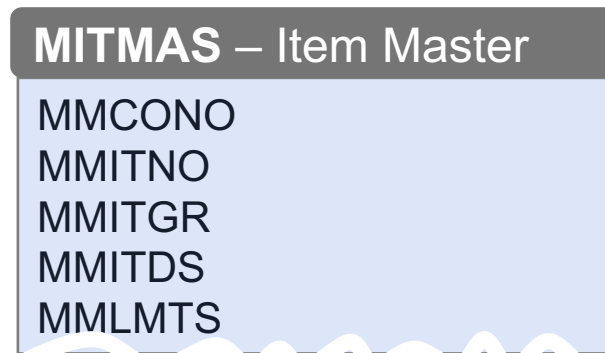
With Compass you don't have to use asynchronous APIs on Data Lake to access data

Main features are:

- Data access using familiar ANSI-SQL standard
- JDBC and API access to Data Lake using Compass
- Data Management and administration in easy-to-use experience in Data Fabric application



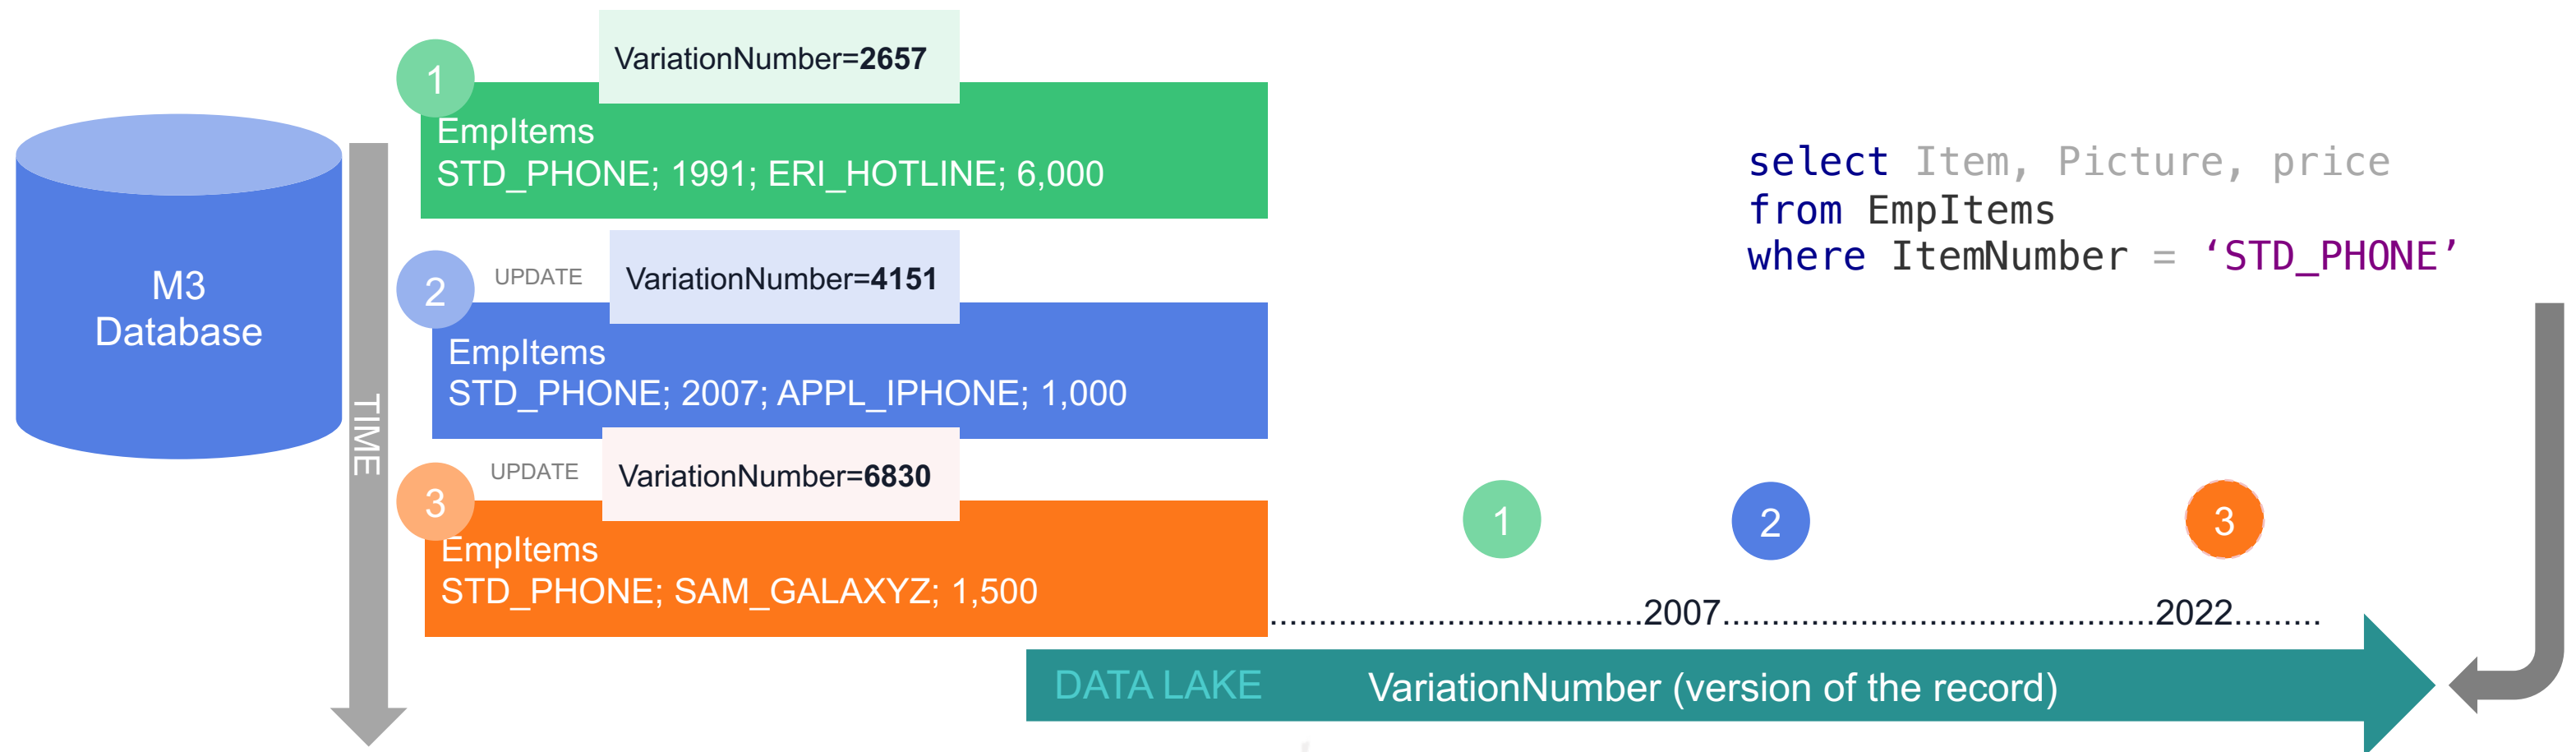
M3 Database tables and Data Lake Data Objects




MMCONO	MMITNO	MMITGR	MMITDS	MMLMTS
100	DRILL001	TOOLS	Drill 3 x P450	2012-12-19T15:33:39.105Z

CONO	ITNO	ITGR	ITDS	LMTS	accounting Entity	variation Number	timestamp	deleted	archived
100	DRILL001	EQX	Drill 3 x P450	2007-11	100_	100102	2007-11	false	false
100	DRILL001	EQT	Drill 3 x P450	2009-12	100_	308134	2009-12	false	false
100	DRILL001	EQT	Drill 3 x P450	2013-03	100_	443001	2013-03	false	false
100	DRILL001	SPP	P450 Classic	2015-01	100_	601834	2015-01	false	false
100	DRILL001	SPP	P450 Classic	2015-12	100_	608808	2015-12	True	false

Variations in Data Lake makes the Timeline



Publish to Data Lake



VariationNumber adds Time to Data

CONO	ITNO	"ITGR"	VariationNumber	Deleted	Timestamp
100	ITEM001	Cheese	1001	-	2004-04-06
100	ITEM001	Dairy	1002	-	2005-05-16
100	ITEM001	Dairy	1038	True	2009-07-31
100	ITEM002	Juice	1003	-	2004-04-06
100	ITEM002	Dairy	1291	-	2005-05-15

A record that is updated or deleted gets a higher variationNumber

SELECT ITNO, ITGR
FROM FAKEMAS



ITEM002 Dairy

= correct

SELECT Sum(Sales)
FROM CustomerOrders
WHERE ITGR = 'Juice'
GROUP BY ITGR...



No result = correct

Let's run 5 Compass Queries on MITMAS (Item Master)

```
SELECT CONO, STAT, ITNO, ITDS, FUDS,  
accountingEntity, deleted, timestamp, variationNumber  
FROM MITMAS WHERE ITNO = 'DRILL001-01'
```

Select the maximum variation excluding deleted and archived records

```
SELECT CONO, STAT, ITNO, ITDS, FUDS,  
accountingEntity, deleted, timestamp, variationNumber  
FROM infor.includedeleted('MITMAS') WHERE ITNO = 'DRILL001-01'
```

Return data, with the highest variation number, that is not archived but could be marked as deleted

```
SELECT CONO, STAT, ITNO, ITDS, FUDS,  
accountingEntity, deleted, timestamp, variationNumber  
FROM infor.allvariations('MITMAS') WHERE ITNO = 'DRILL001-01'
```

Return all variation numbers of data, also deleted but not archived

```
SELECT CONO, STAT, ITNO, ITDS, FUDS,  
accountingEntity, deleted, timestamp, variationNumber  
FROM infor.ShowArchived('MITMAS') WHERE ITNO = 'DRILL001-01'
```

Return only all archived data with highest variation numbers and also deleted

39

```
SELECT CONO, STAT, ITNO, ITDS, FUDS, accountingEntity, deleted, archived, timestamp, variationNumber  
FROM infor.AllVariationsIncludeArchived('MITMAS')
```

```
SELECT CONO, STAT, ITNO, ITDS, FUDS, accountingEntity, deleted, archived, timestamp, variationNumber  
FROM infor.ShowAllArchived('MITMAS')
```

Return all data with highest variation numbers plus the archived data

ETL-Client

Extract and transform data from Data Lake and load it into your on-prem relational database

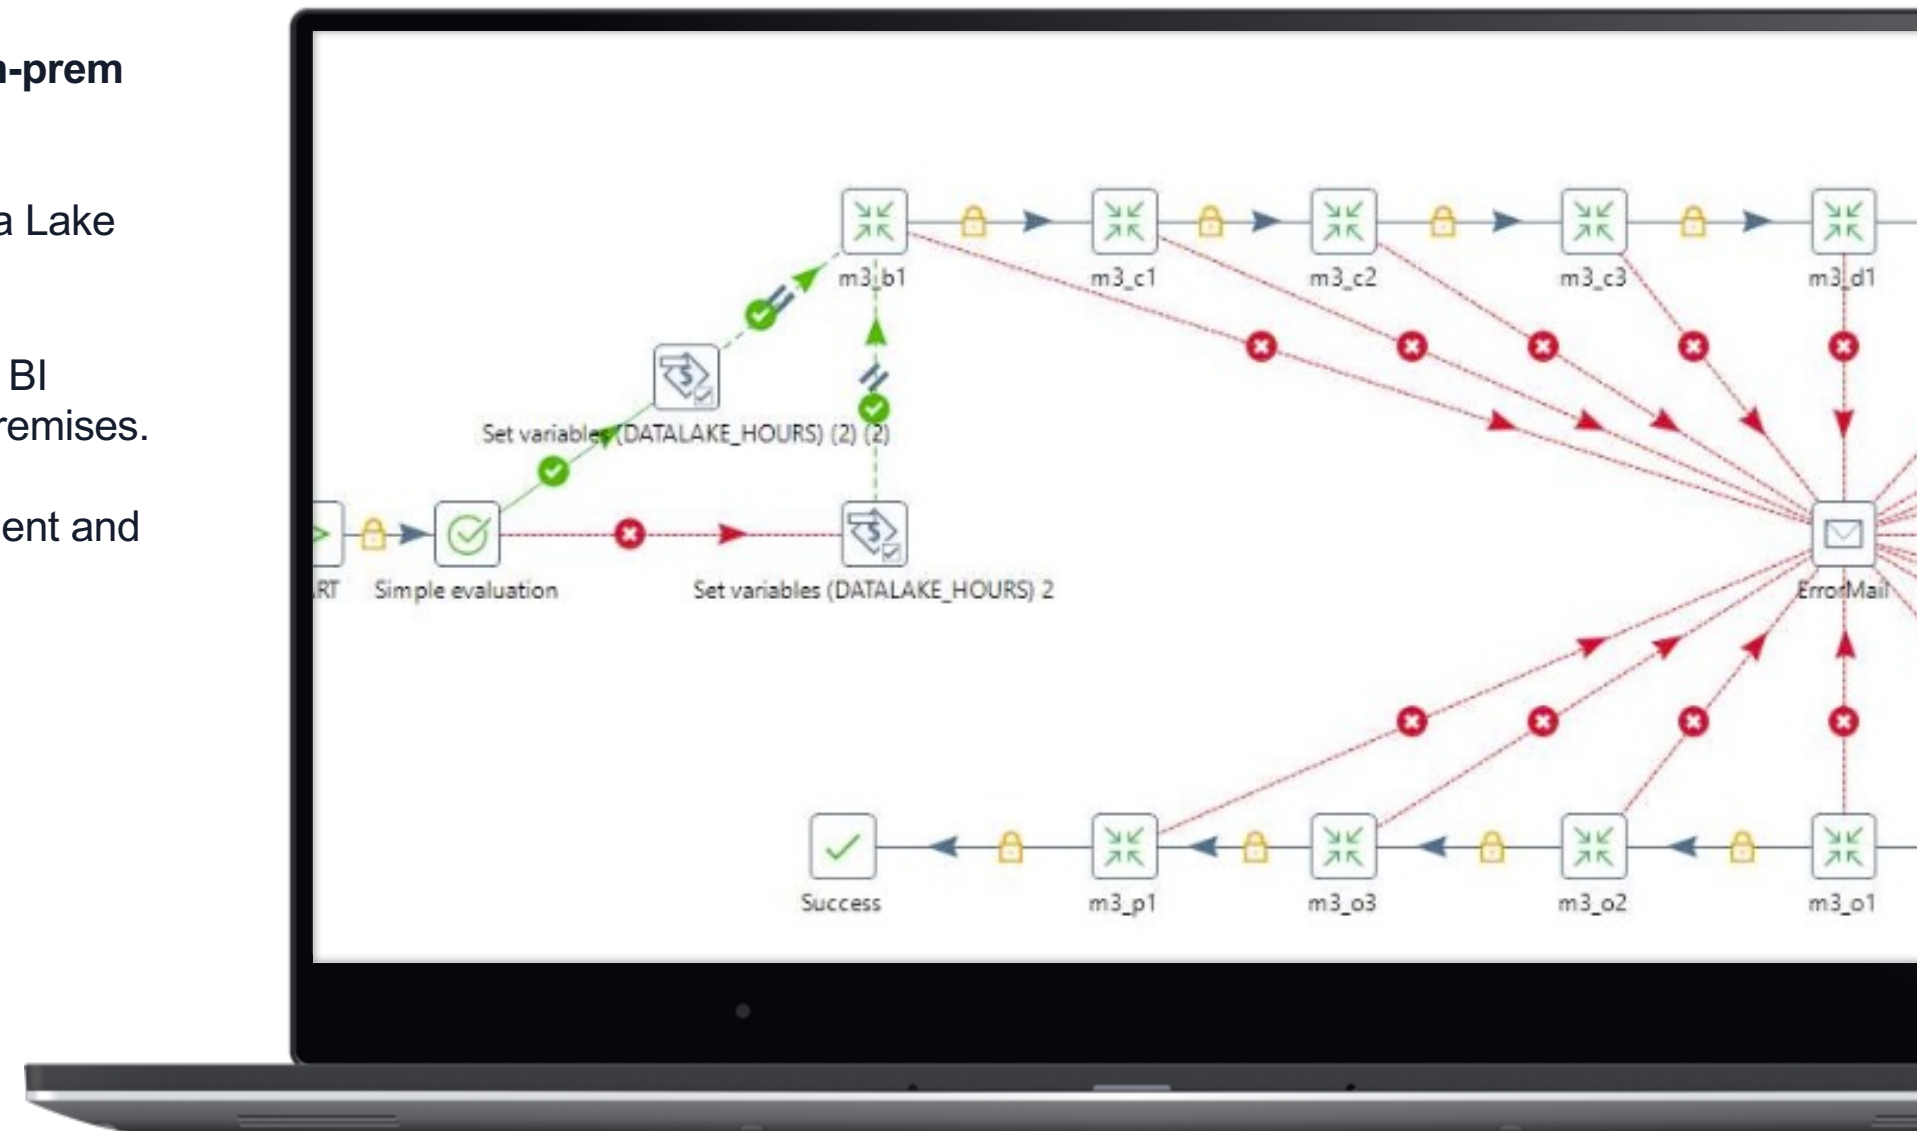
The ETL-Client supports incremental extraction of data ingested to Data Lake and loads it into your database.

This enables a database representation of M3 that can let your existing BI platforms and applications access the data in a similar way as M3 on-premises.

You install and manage the ETL-Client on your own premises environment and incrementally extract data from your Infor CloudSuite Data Lake.

Main features are:

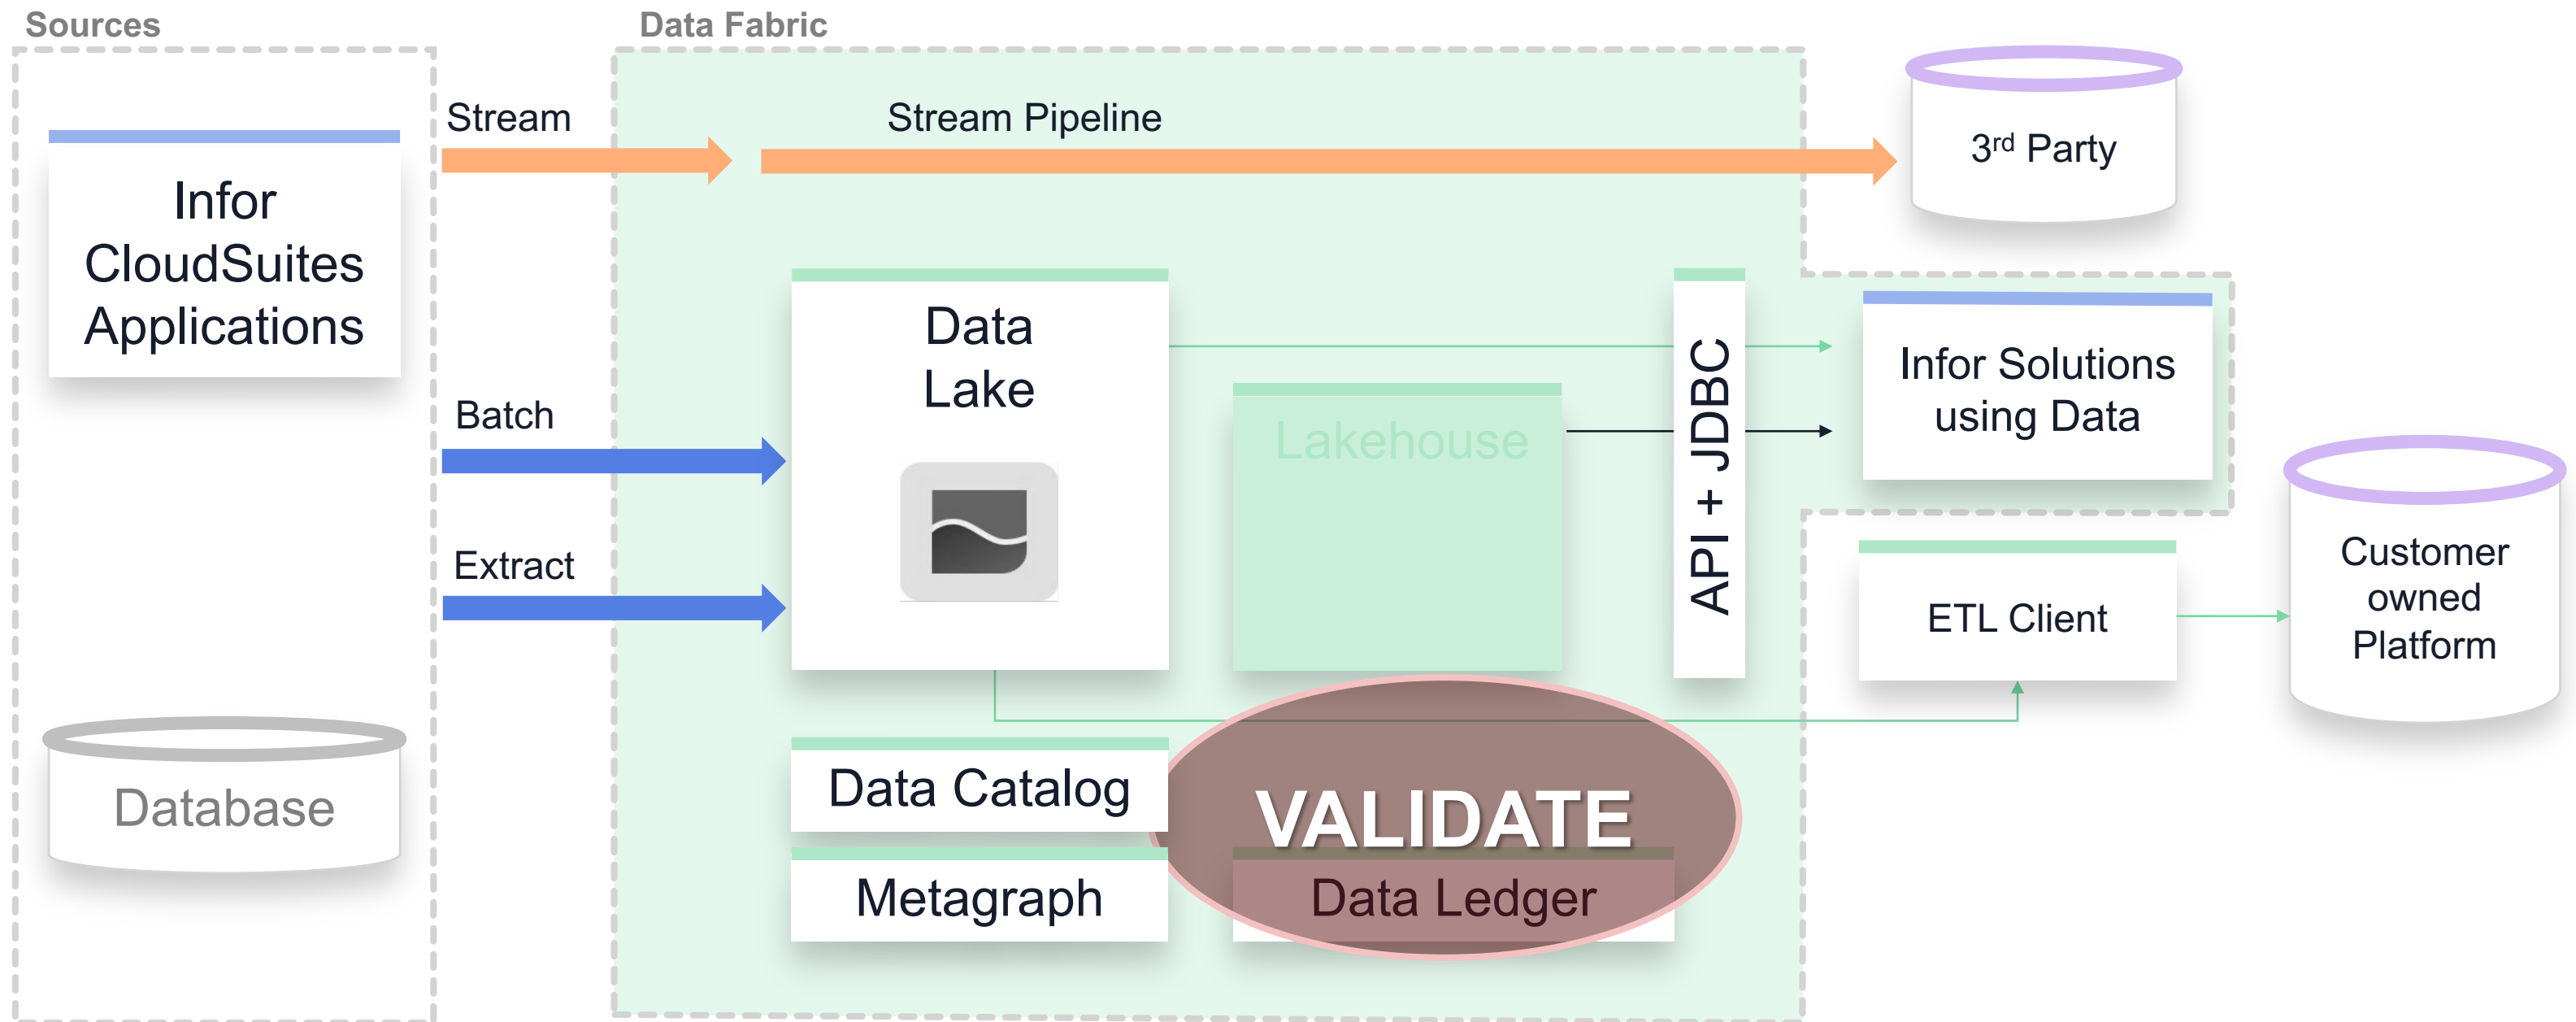
- Scheduled periodic incremental data extraction
- Functionality to transform data before loading into the database.
- Supports several database platforms



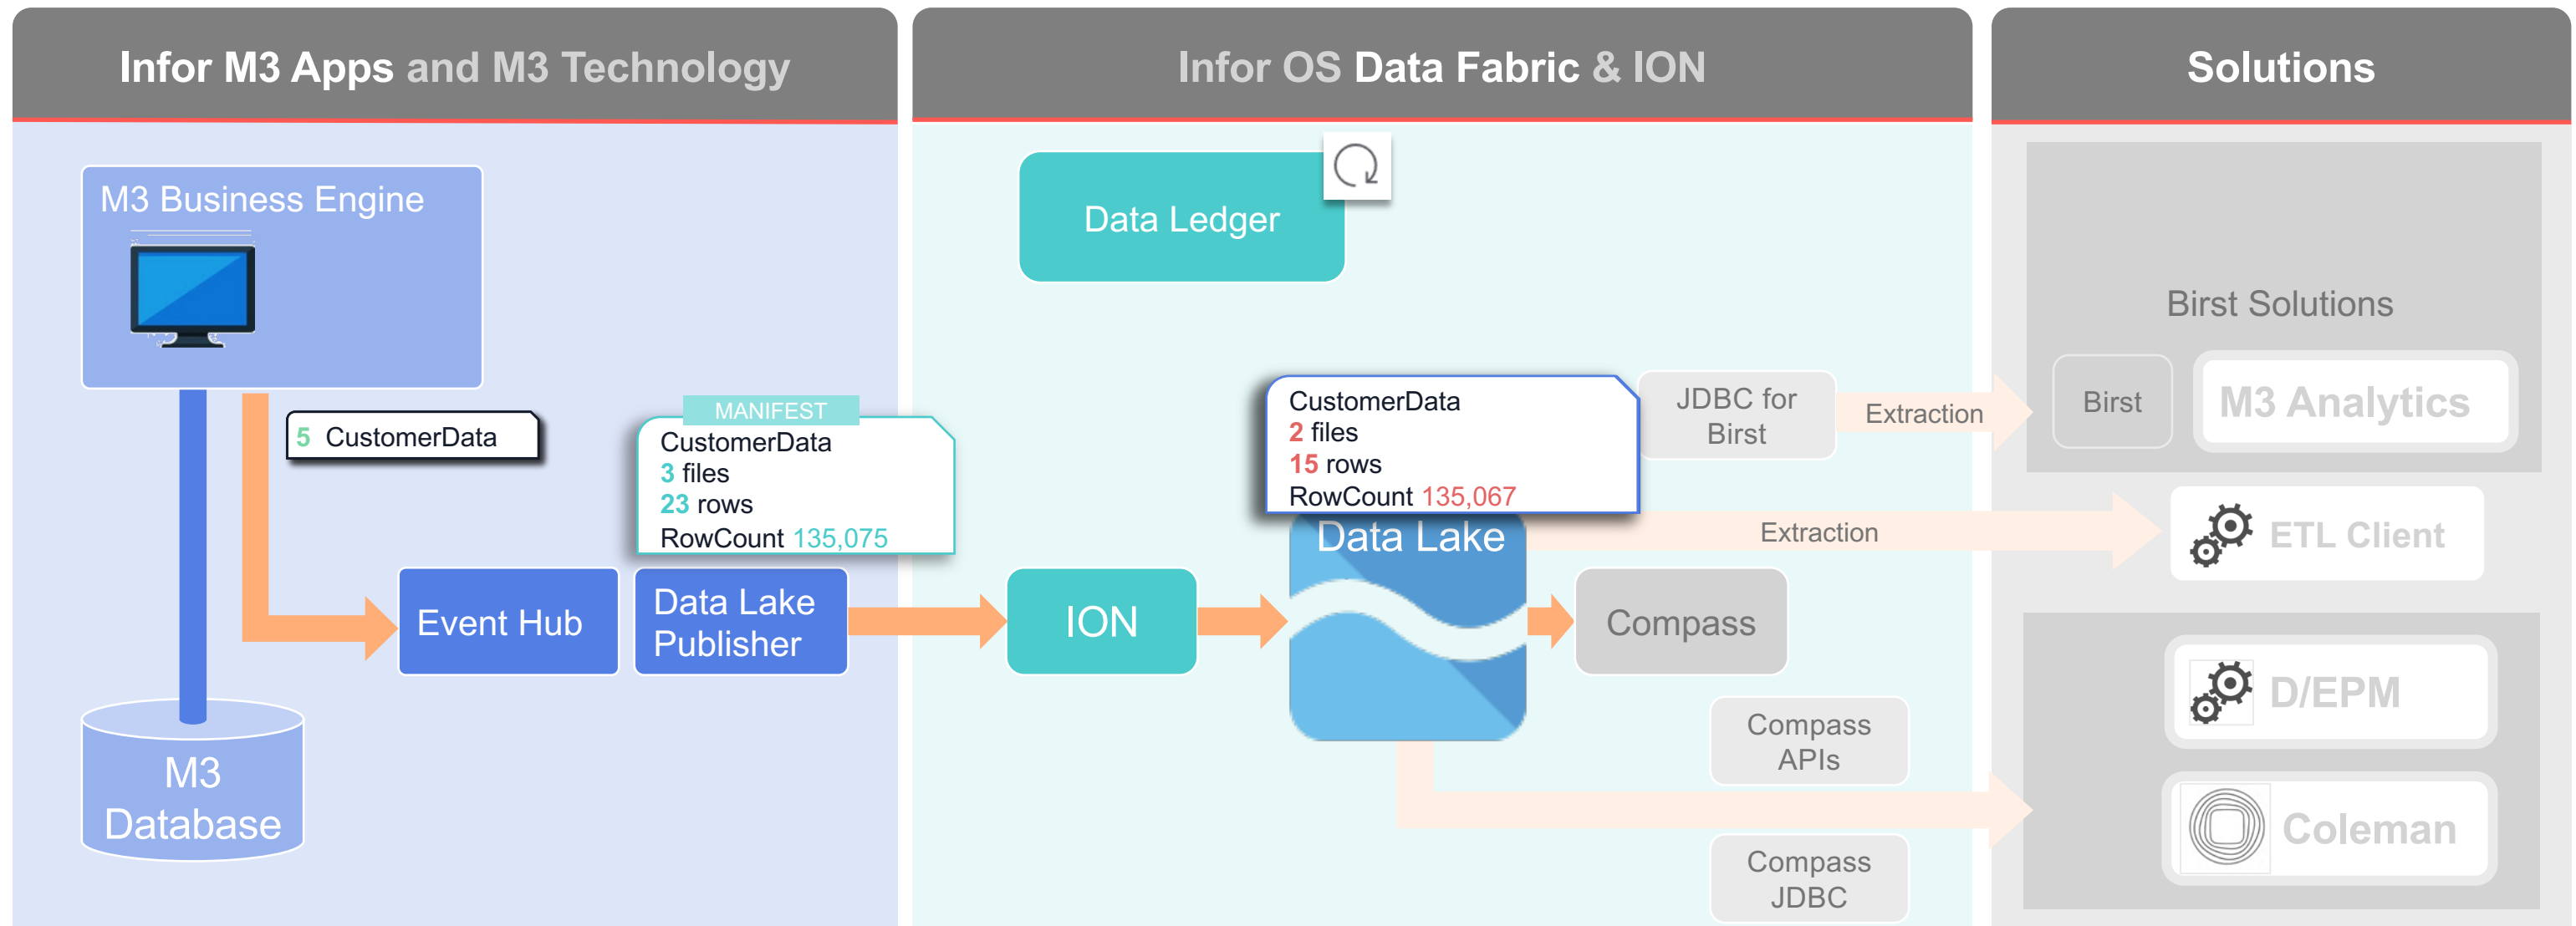


07 Validating and Managing Data Lake

Data Fabric - Validate



M3 and Data Ledger (available 2022.06)



M3 and Data Ledger

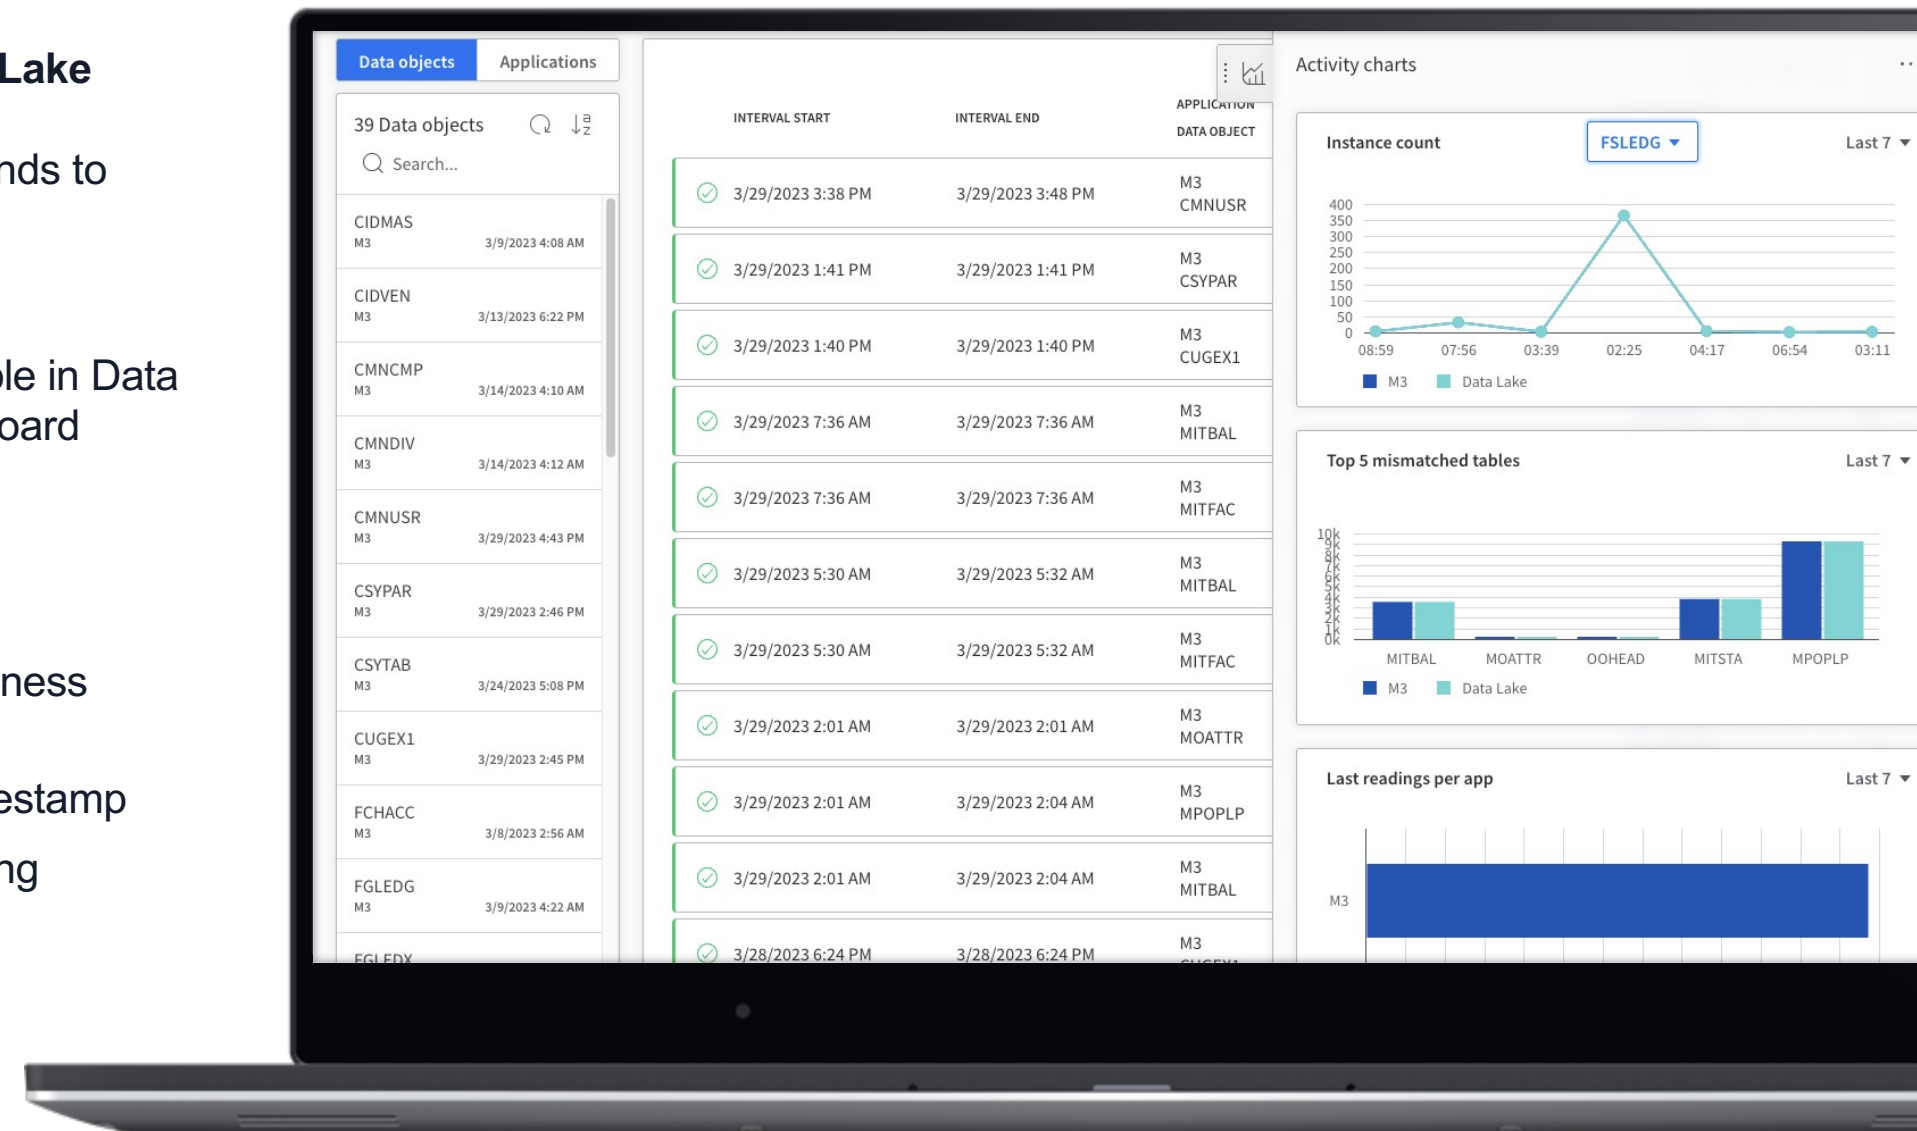
Data Ledger monitors the M3 data publishing and storage in Data Lake

To reconcile Data Lake with M3, Data Ledger is monitoring what M3 sends to Data Lake and what is being stored.

This is the first step in monitoring whether data is published and available in Data Lake in an asynchronous architecture as with ION. A second tab dashboard shows streamed data into Data Fabric stored in Data Lake.

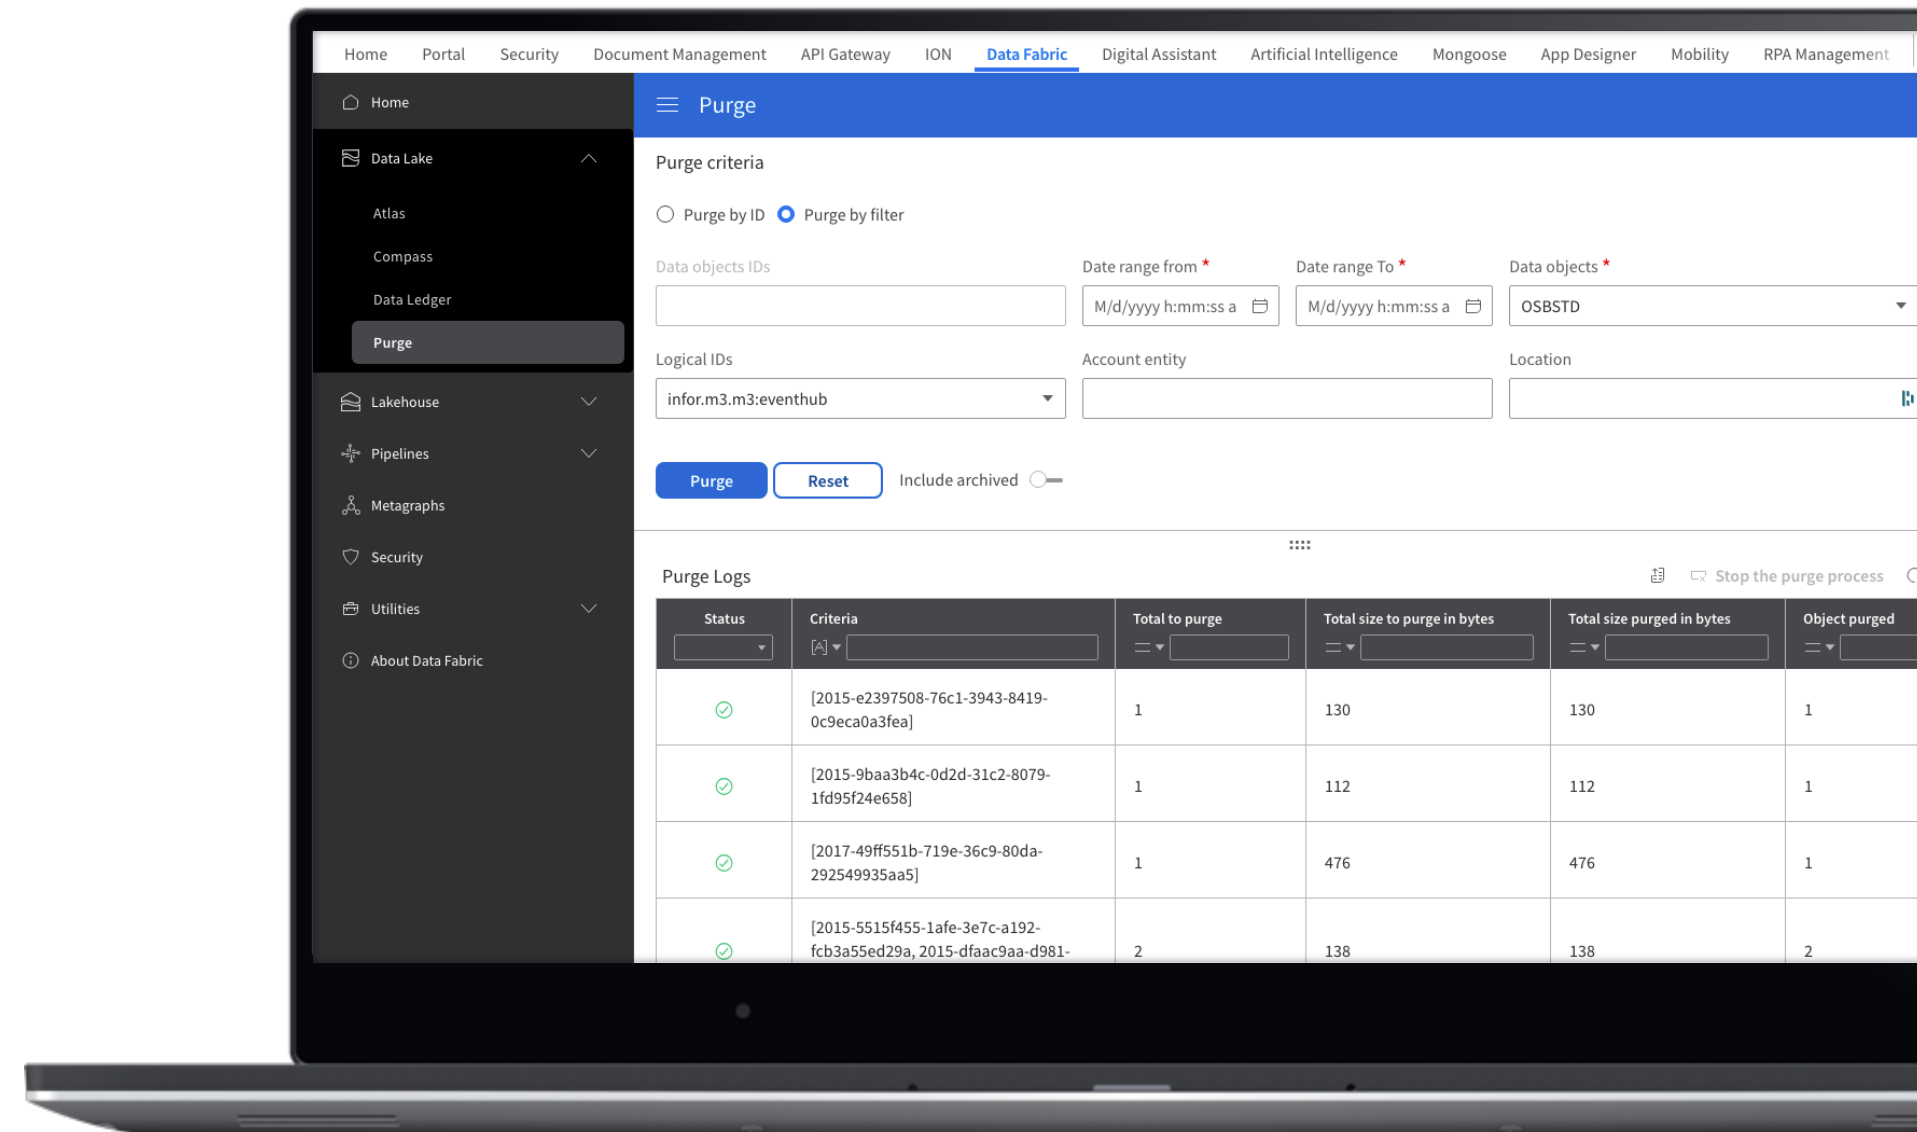
Main features are:


- M3 sends periodically information about table row count in M3 Business Engine Database as of a timestamp
- Data Ledger looks into Data Lake for row count as of the same timestamp
- Statistics presented in Data Ledger with red or green cards indicating successful or deferred or failed data delivery



Deleting data from Data Lake

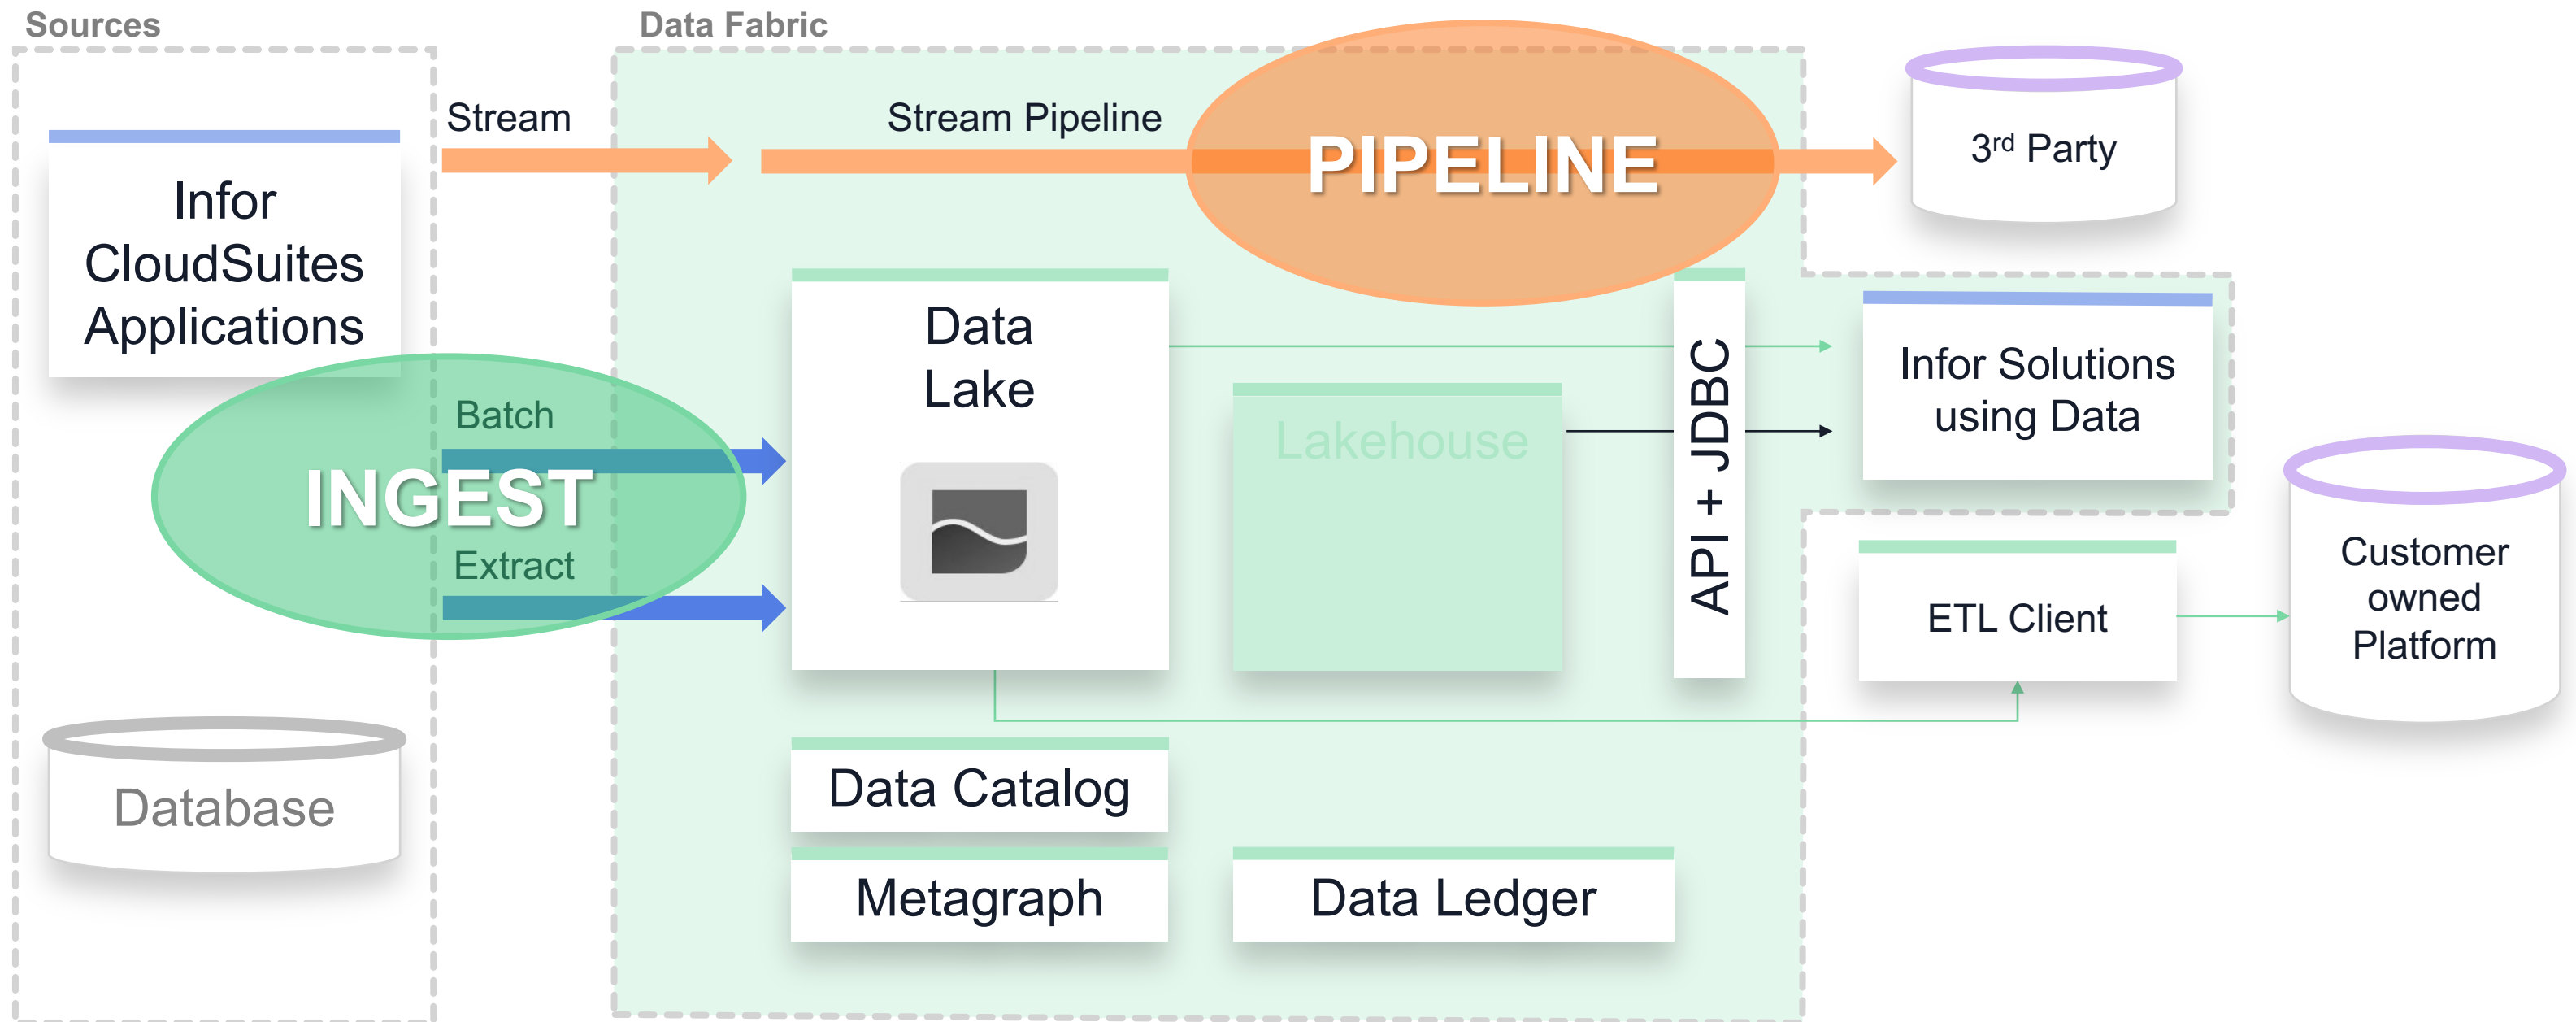
- Purge
- Mark Corrupt
- Clear Compass cache





08 M3 Streaming & Stream Pipelines

Data Fabric – Streaming & Stream Pipelines



Data Fabric – Streaming & Stream Pipelines

The screenshot displays the Data Fabric interface for a stream pipeline named 'MITMAS_UPSERT'. The left sidebar contains navigation options: Home, Data Lake, Pipelines (with sub-items: Stream Pipelines, Destinations, Replay Queue), Metagraphs, Security, Utilities, and About Data Fabric. The top navigation bar shows the pipeline name and tabs for 'Model' and 'Overview'. The 'Stream' status is 'Running'. Other controls include 'Initial Load', 'Replay queue', and various utility icons. The main area shows a diagram of the pipeline with two components: 'M3 MITMAS' (Subscription) and 'MITMAS Current da...' (Delivery). Below the diagram, there are two colored bars representing the components: a purple bar for 'Subscription' and an orange bar for 'Delivery'.

Destinations

Pipelines enables fast data delivery to various technologies, relational databases, analytics warehouses, streaming platforms and storage locations.

The Destinations component is used for defining and managing the connection to these locations where Stream Pipelines can offload data in real-time processing.



Amazon Aurora PostgreSQL

(supported in first release)



Azure Database for PostgreSQL



Data Fabric Lakehouse

(2H 2024)




Snowflake

(1H 2024)



SQL Server

(2H 2024)



Amazon Redshift



Streaming platforms



Storage

09

News in 2023.10 and future roadmap



Continuous Improvement

Compass Performance & Cleansing

Compass SQL sees substantial query performance tuning improvements including functional updates to further cleanse & de-duplicate content.

Value

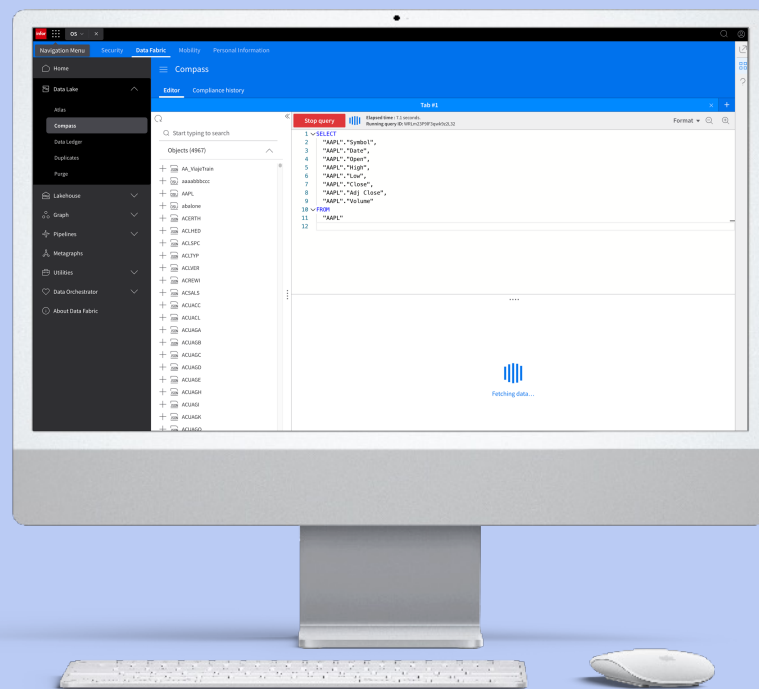
- Improved query performance
- Reduced caching wait times
- Accurate & effective data presentation aligned with functional business users' expectations
- Data de-duplication



The screenshot displays the Infor Compass SQL interface. The top navigation bar includes Home, Portal, Security, Data Fabric, Mobility, and Personal Information. The main interface is divided into three sections: a left sidebar with navigation options like Home, Data Lake, Atlas, Compass, Data Ledger, Duplicates, Purge, Lakehouse, Graph, Pipelines, Metagraphs, Utilities, Data Orchestrator, and About Data Fabric; a central pane showing a search bar and a list of objects (4826) including NetbackandOrderHistory, Nexus_OrderItemPOC2, NG_CostCenterLoad, nk_bigint_to_timestamp, notes, num_0, OAGRHE, OAGRNL, OAGRTP, objectname_sk, OCUSAD, and OCUSMA; and a right pane with a 'Run query' editor and a results table. The query editor shows a SQL statement: 'CREATE VIEW "CustomerOrders" AS SELECT "OCUSMA"."accountingEntity", "OCUSMA"."variationNumber", "OCUSMA"."timestamp", "OCUSMA"."deleted", "OCUSMA"."archived", "OCUSMA"."CONO", "OCUSMA"."DIVI", "OCUSMA"."STAT", "OCUSMA"."CUNO", "OCUSMA"."CUCL", "OCUSMA"."CUTP", "OCUSMA"."ALCU", "OCUSMA"."CUNM", "OCUSMA"."CUA1", "OCUSMA"."CUA2", "OCUSMA"."CUA3", "OCUSMA"."CUA4",'. The results table shows 47 rows with columns: CONO, DIVI, STAT, CUNO, CUCL, CUTP, ALCU, CUNM, and CUA1. The table contains data for various customer types and addresses.

	CONO	DIVI	STAT	CUNO	CUCL	CUTP	ALCU	CUNM	CUA1
1	750		20	Y10001	Y06	0	TRADE	Trade account customer	Address Line 1
2	750		20	Y001000001	ZZZ	0	SHOP	Generic Shop Customer	Street/P.O. Box
3	750		20	Y001000000	ZZZ	0	SHOP	Generic Shop Customer	Street/P.O. Box
4	750		20	Y001000003	ZZZ	0	SHOP	Generic Shop Customer	Street/P.O. Box
5	750		20	Y10003	Y02	2	RETAIL	Large Retailer	Address Line 1
6	750		20	Y10000	Y01	0	DOMESTIC	Local customer	Local Customer Address Line 1
7	750		20	Y10003CA	Y02	0	RETAIL CA	Large Retailer California Branch	Address Line 1
8	750		20	Y900000008	ZZZ	8	TEMPORARY	Lombardo Incorporated	2356 Sea Shore Drive
9	750		20	Y10003IL	Y02	0	RETAIL IL	Large Retailer Illinois Branch	Address Line 1

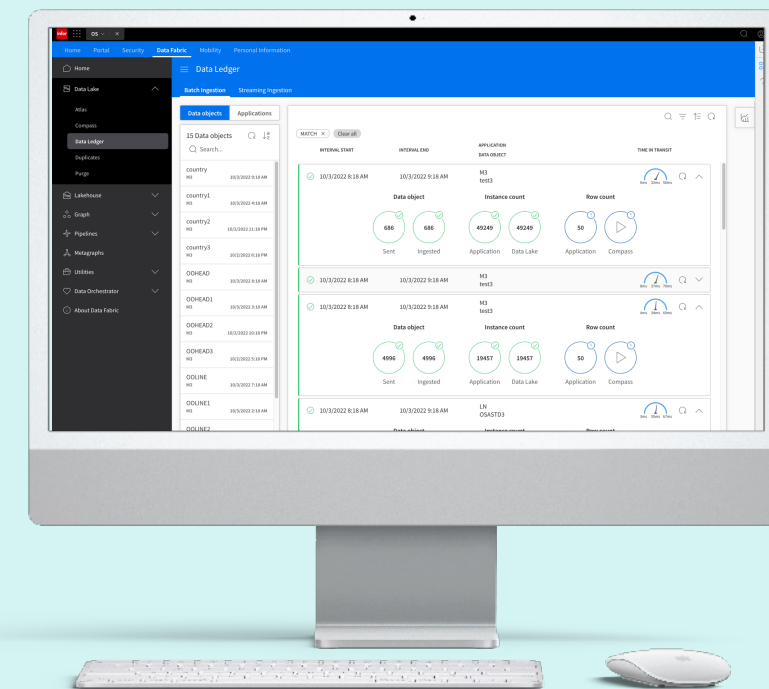
Performance



Compass SQL Speed

- Across-the-board query performance boost
- Effectively eliminate “caching” penalties felt by customers during queries
- More effectively troubleshoot Compass SQL with stored procedures
- Faster, more flexible troubleshooting based on shortened rebuild times
- Feature opportunities to expose custom data partitioning for additional performance gains

Data Quality



Compass SQL Accuracy

- Accurate reflection of data consistent with ERP application data
- Additional data de-duplication measures for greater accuracy & quality
- Eliminate customer-implemented workarounds that otherwise reduce query performance or incur additional Compass compute time

Access Entitlements

Fabric Security Framework

An expanding selection of query tools, data browsing experiences, and streaming capabilities requires segregation of duties. The security framework for Data Fabric allows you to implement principles of least privilege to entitle access to the big data tools users need.

April 2023

Value

- IFS Security Role integration
- Two default IFS roles & unlimited custom roles for ACL design
- Principle of least privilege (entitle up)
- User entitlements perspective



The screenshot displays the 'Access Controls' interface in Microsoft Fabric. On the left, a navigation pane includes 'Home', 'Data Lake', 'Data Pipelines', 'Security', 'Downloads', and 'About Data Fabric'. The 'Security' section is expanded to show 'Access Controls'. The main content area is titled 'Access Controls' and features a 'Users (20)' list with a search bar and a 'Sort by Name' dropdown menu (set to 'Ascending'). The user list includes Julie Holland (Product Manager), Catherine Dixon (Business Analyst), Douglas Sanders (Business Analyst), Catherine Torres (Manager), Frank Day (Data Analyst), and Douglas Sanders (Business Analyst). The right side of the interface shows a 'Permissions' section for 'Security Roles' with a user profile for Catherine Dixon (Business Analyst). Below this, there are sections for 'Data Lake' and 'Data Pipelines'. The 'Data Lake' section includes permissions for 'Atlas', 'Compass', 'Ledger', and 'Purge', with checkboxes for actions like 'Download', 'Suppress', 'Purge', 'Query', 'Manage views', 'Run administrative stored procedures', 'Export', 'Purge by ID', 'Purge by filter', and 'Export'. The 'Data Pipelines' section includes permissions for 'Destinations' and 'Stream Pipelines', with checkboxes for 'Manage', 'Import', and 'Export'.

Data Entitlements

Compass Security

Object-level and column-level data security rules entitle and restrict access to content in storage and data accessible through Compass SQL based on Infor Federated Security security roles.

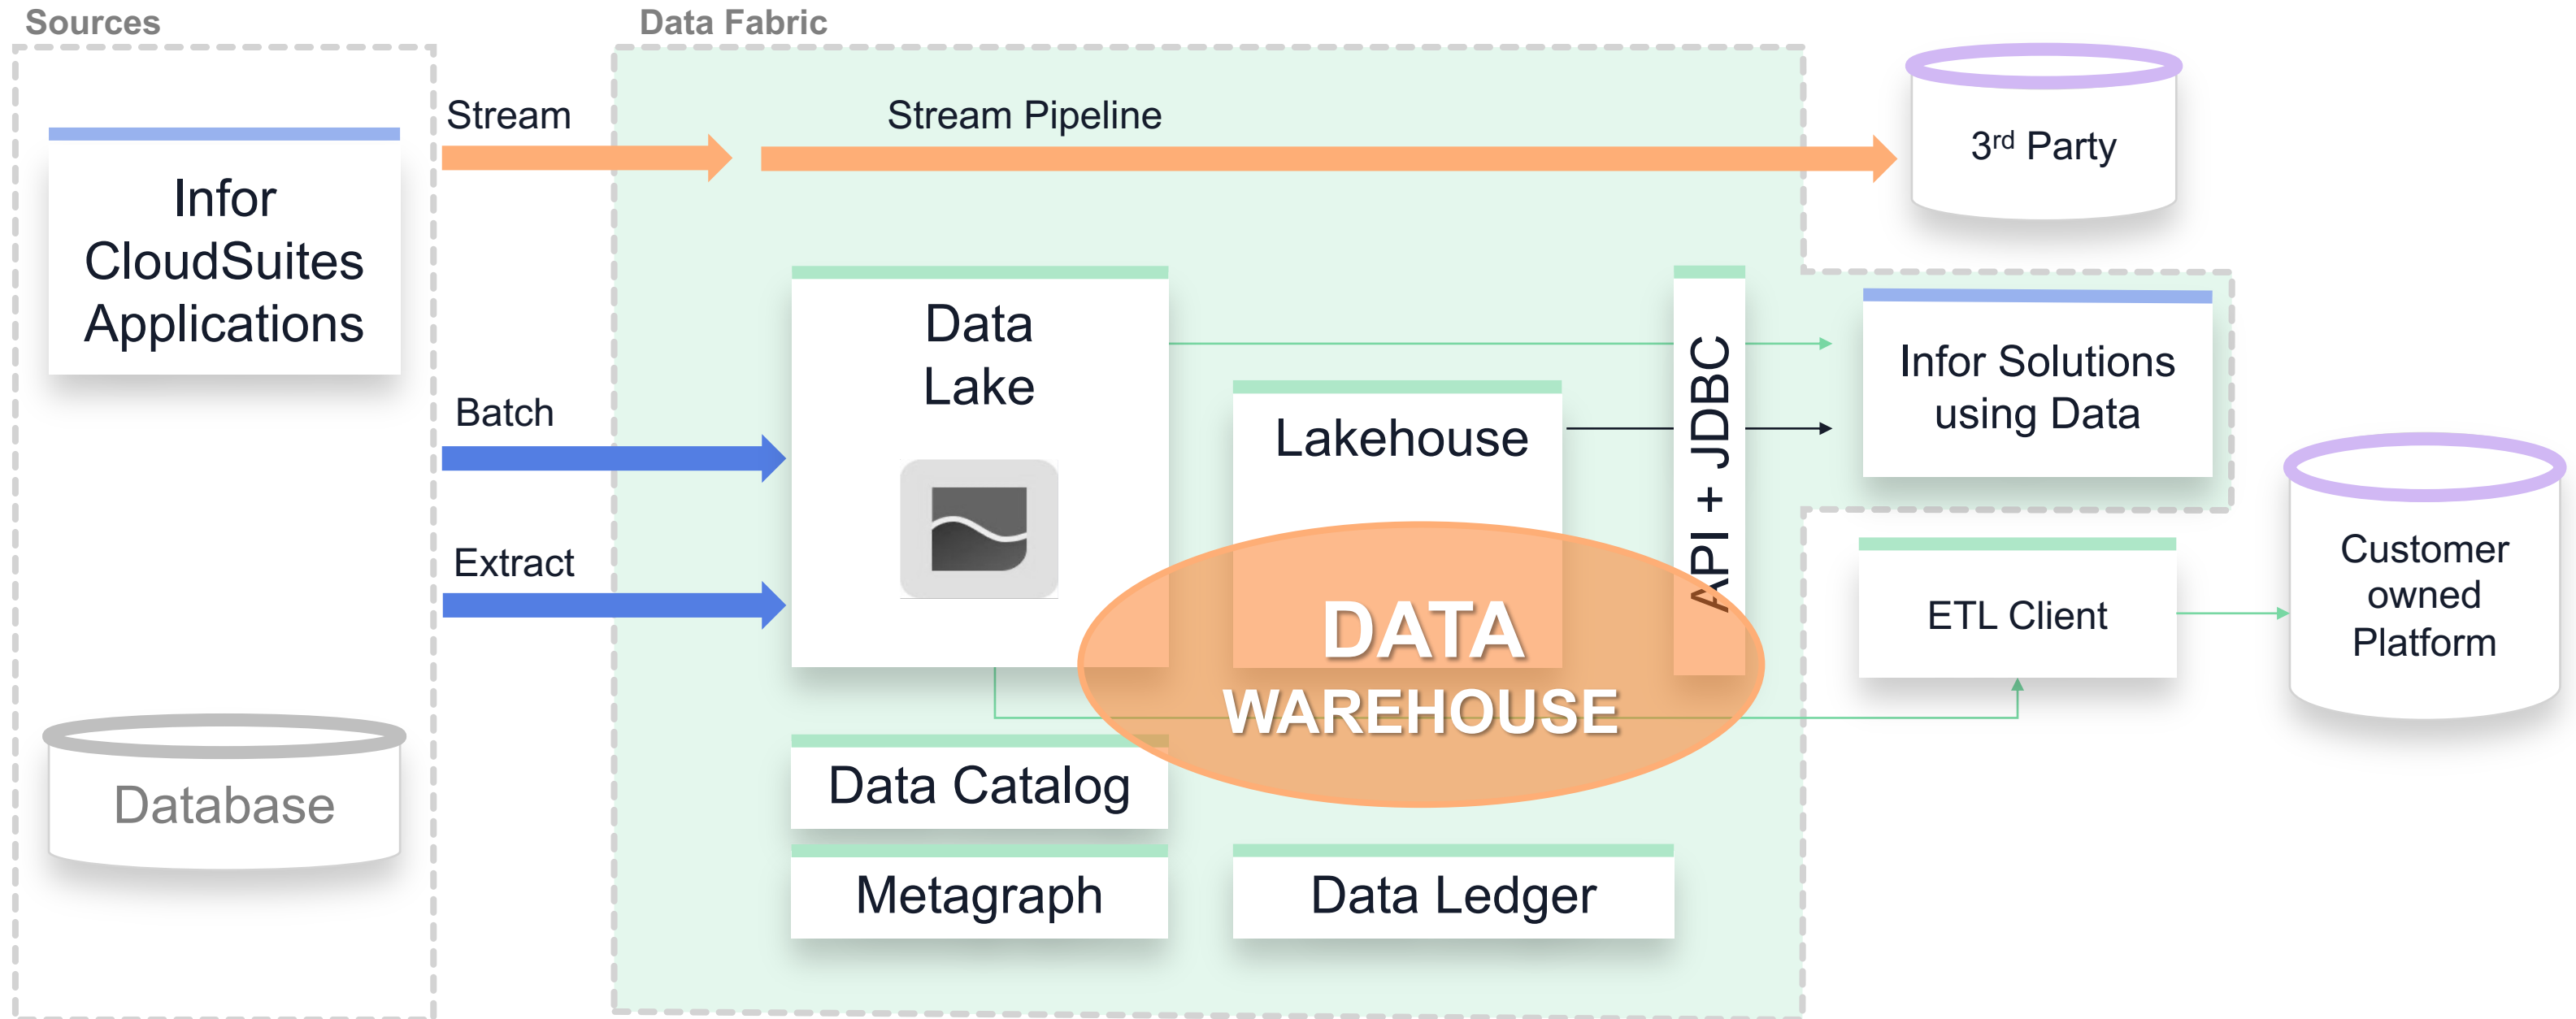


The screenshot displays the Compass Security interface. On the left is a dark sidebar with navigation options: Home, Data Lake, Lakehouse, Graph, Pipelines, Metagraphs, Security (highlighted), Utilities, and About Data Fabric. The main content area is titled 'Access Controls' and has tabs for 'Security Roles' and 'Users'. The 'Users' tab is active, showing a list of 50 users. The user 'Catherine Dixon' (Business Analyst) is selected. To the right, the 'Data Permissions' section shows a table of permissions for Catherine Dixon. The table has columns for 'OBJECT NAME', 'OBJECT TYPE', and a dropdown for access level. Below this, a table shows 'ACHLED' permissions with columns for 'Property Name', 'Title', and 'Description'.

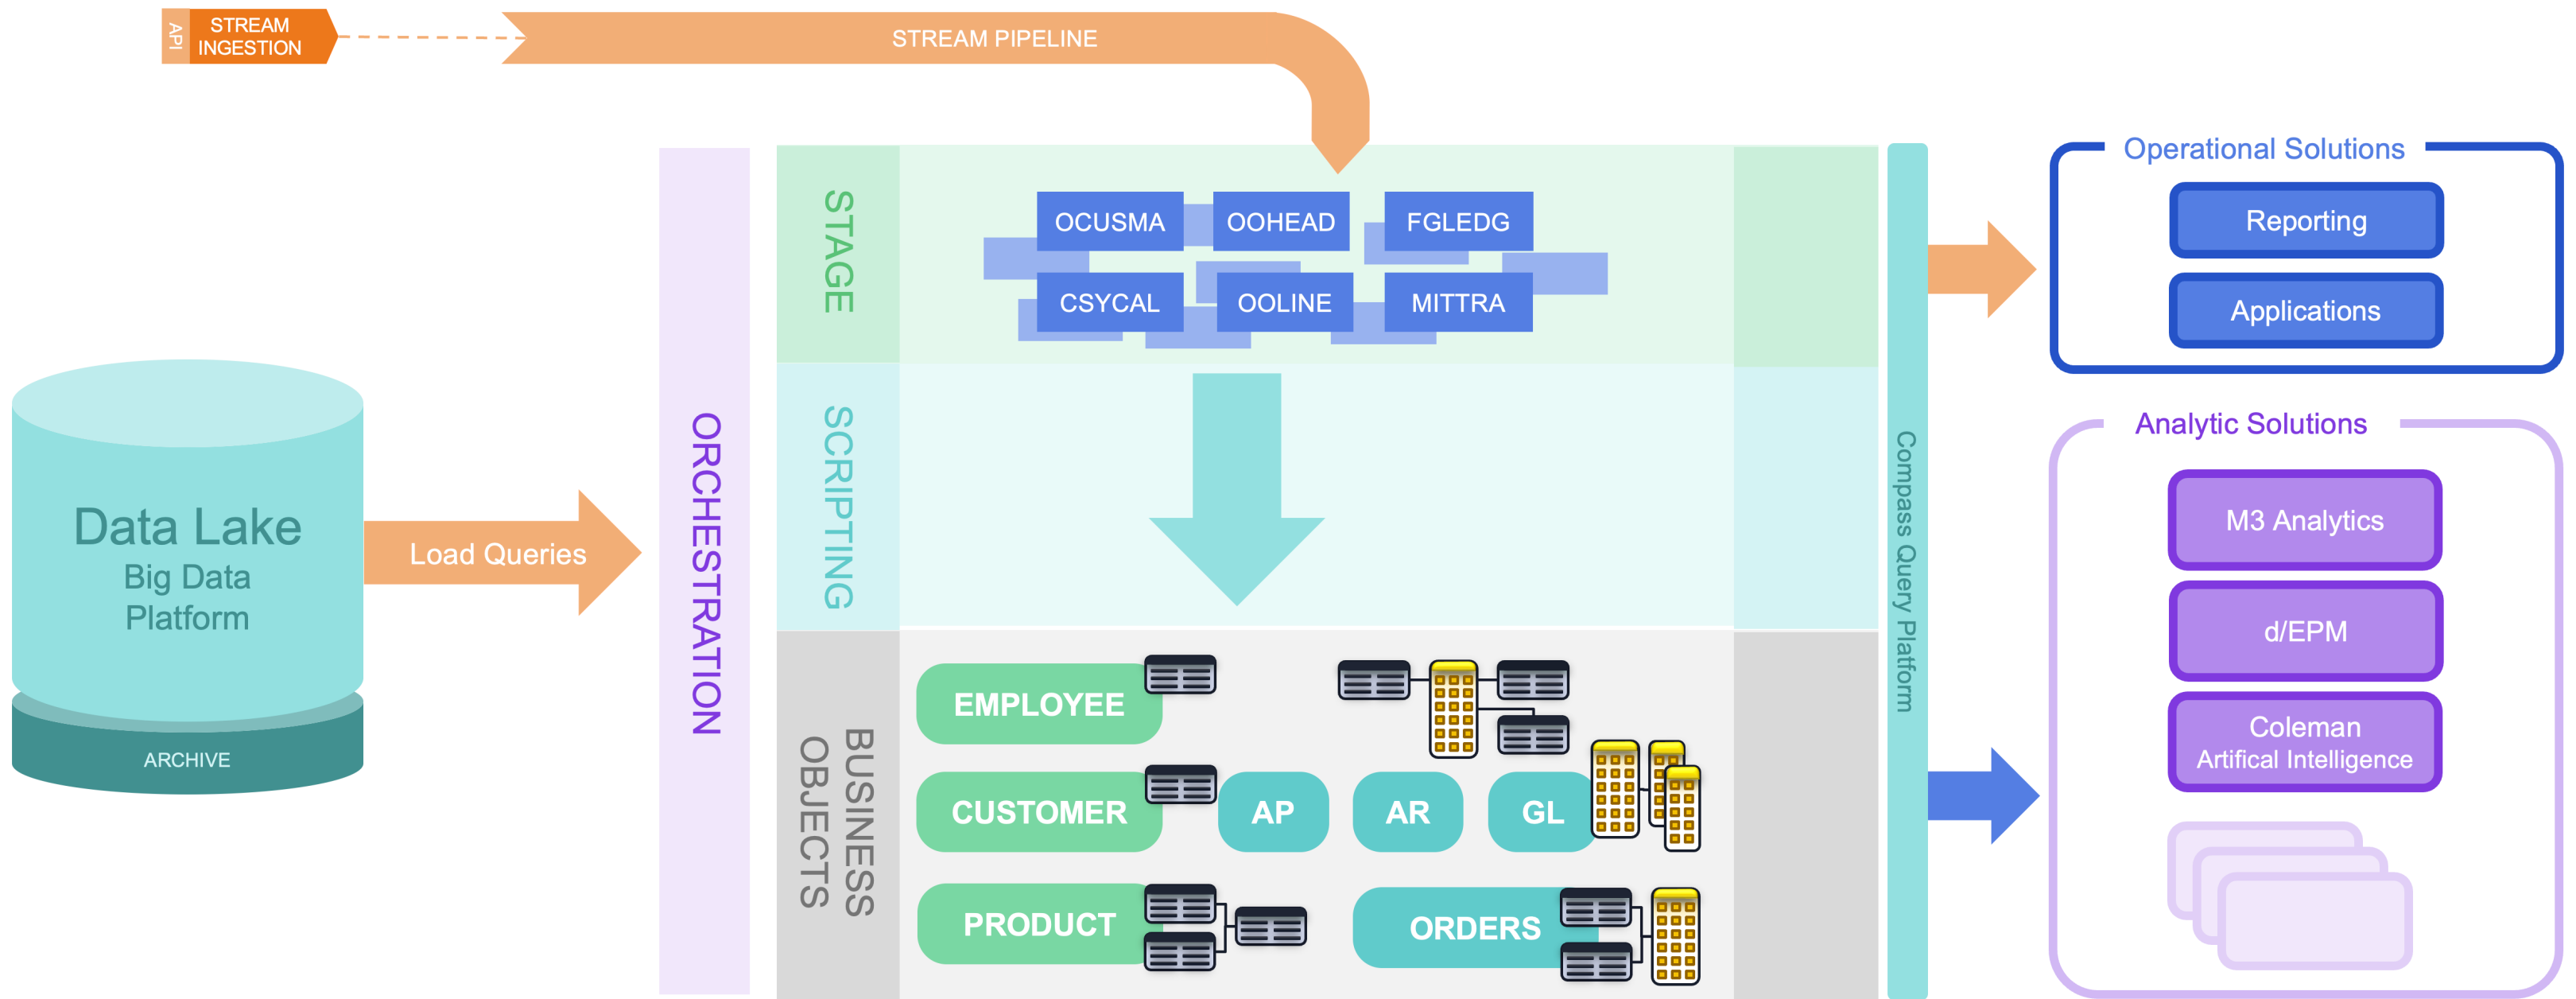
OBJECT NAME	OBJECT TYPE	Access Level
AANAME	JSON	Full access
AnyDocument	ANY	Full access
BB_Customer	DSV	Full access
BB_Payment	DSV	Full access
BB_Payment	DSV	Full access
Customers	JSON	Partial access
ACHLED	JSON	Partial access

Property Name	Title	Description	
<input checked="" type="checkbox"/>	Lorem	Ipsum	Lorem ipsum dolor sit amet, consectetur.
<input checked="" type="checkbox"/>	Lorem	Ipsum	Lorem ipsum dolor sit amet, consectetur.
<input checked="" type="checkbox"/>	Lorem	Ipsum	Lorem ipsum dolor sit amet, consectetur.
<input checked="" type="checkbox"/>	Lorem	Ipsum	Lorem ipsum dolor sit amet, consectetur.
<input type="checkbox"/>	Lorem	Ipsum	Lorem ipsum dolor sit amet, consectetur.
<input checked="" type="checkbox"/>	Lorem	Ipsum	Lorem ipsum dolor sit amet, consectetur.
<input checked="" type="checkbox"/>	Lorem	Ipsum	Lorem ipsum dolor sit amet, consectetur.

Data Fabric – Streaming & Stream Pipelines



M3 and Data Fabric – Lakehouse and Real Time Data



Data Lake vs Data Warehouse

Data Lake



Raw data files from data sources
Metadata defined in Data Catalog

- ✓ Stores raw, unstructured, or semi-structured data
- ✓ Suitable for storing vast amounts of diverse data types
- ✓ Supports data exploration, analytics, and machine learning
- ✓ Schema-on-read approach, meaning the structure is applied when data is read
- ✓ More flexible in accommodating changes in data sources and formats
- ✓ Generally used for long-term storage and big data analysis


Operational Data Store



An Operational Data Store (ODS) serves as a near real-time repository that consolidates and stores current operational data from multiple sources

- ✓ Stores processed, structured, and current operational data
- ✓ Optimized for real-time or near-real-time access
- ✓ Supports operational reporting, decision-making, and processes
- ✓ Data is structured
- ✓ Typically integrates data from multiple sources to provide a unified view
- ✓ Focuses on providing quick access to data for operational tasks

Data Warehouse (Lakehouse*)



Suite of industry data models with predefined data enrichment

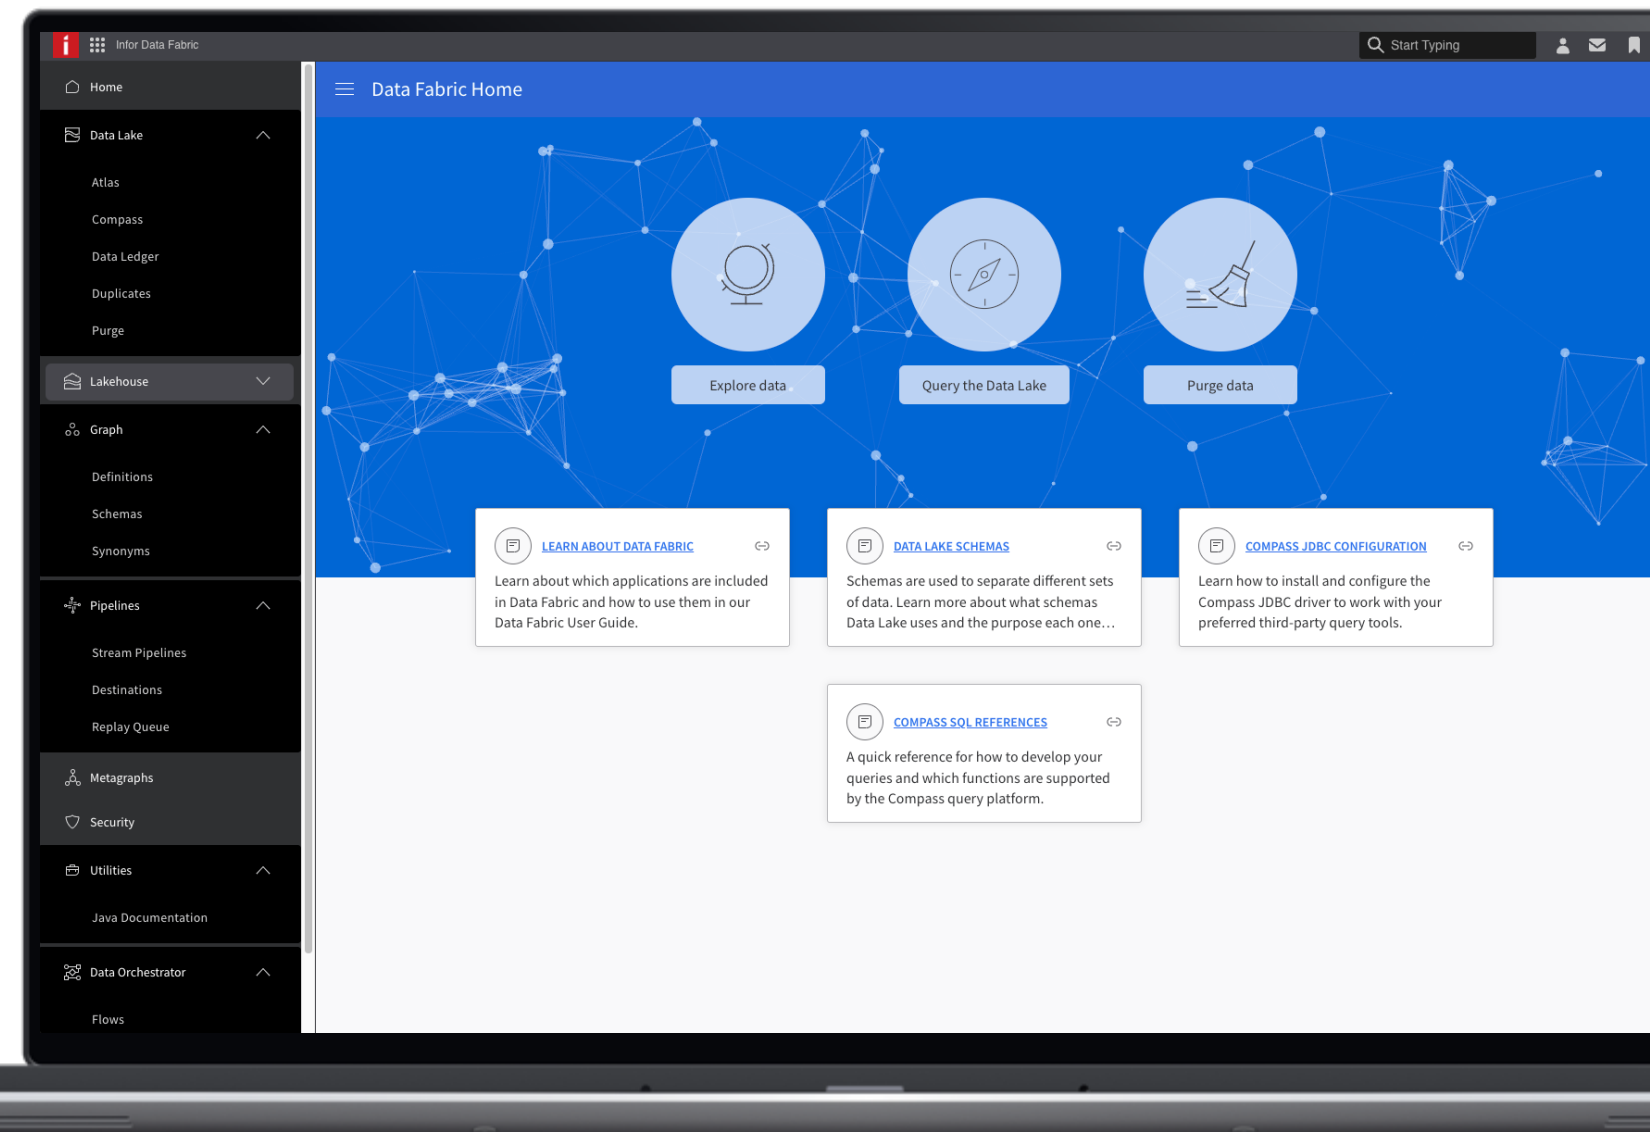
- ✓ Fully governed, curated data models
- ✓ Schema on-write – conforming data to well-defined target structures
- ✓ Business intelligence & operational reporting consumers
- ✓ Structured data in relational systems

* Planned General Available April 2024

M3 and Data Fabric - Summary

M3 has chosen to collaborate closely with Infor OS Data Fabric

- M3 started early to adopt **Data Lake** back in 2018 as the data source for M3 Analytics, using Birst and JDBC against Data Lake. M3 sent data in micro-batches to an ION Data Lake Flow for storage in Data Lake.
- **Compass** on top of Data Lake made customers start querying Data Lake with standard SQL and integrated their downstream applications to work similarly to how SQL Queries were used in the on-prem platform.
- When business-critical applications were using Data Lake data, it was necessary to reconcile data in Data Lake, and **Data Ledger** was developed by the Data Fabric team.
- To make Data Fabric support direct ingestion without ION, **batch** ingestion, and **stream ingestion** methods were built. M3 is the first Infor solution to use streaming and several M3 solutions are using the batch ingestion method.
- Several customers had on-prem SQL Servers for their BI platforms and the **ETL Tool** was added to Data Fabric to support incremental ETL from Data Lake.
- To support operational data use cases with real-time requirements, **Data Fabric Stream Pipelines** feature was built and M3 can now stream data in real-time to an AWS Aurora PostgreSQL or Azure PostgreSQL database.
- **Lakehouse** is the next area for M3 to adopt and build solutions for Infor's Data Warehouse as a Service for open **Decision Support** solutions and **operational** real-time data access solutions.



10 Licensing and Info



Licensing add-on Stream Pipelines

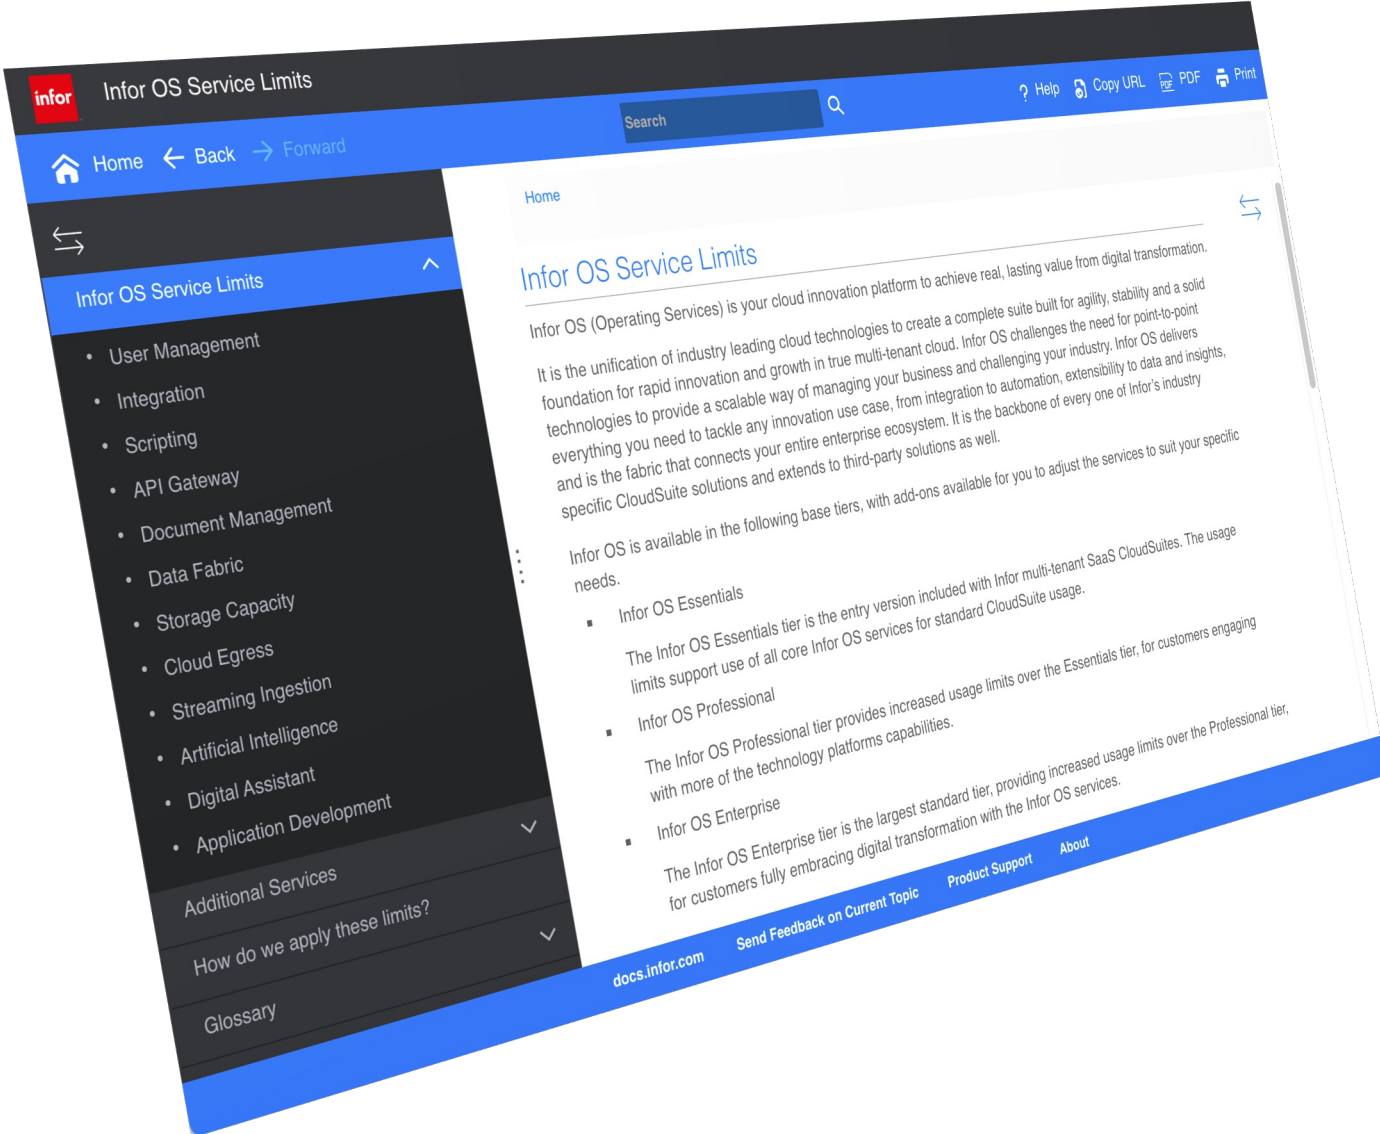
[Additional Services - Stream Pipelines \(ION-S-DFP\)](#)

How to increase your service limits

Additional subscription quantities can be obtained using add-on SKU:	ION-S-DFP
Add-on SKU sold with license method (LM) of:	TECH
Add-on quantities available in multiples of:	165,000
Add-on increments service resource:	Events Processed

Service Usage Limits

Resource	Type	Subscription Quantity 1 to n Usage Entitlement	TECH License Method Unit	Measured @ Tenant/ Customer
Events	Add-on	Subscription quantity	Events per day	Customer



Data Fabric Cheat Sheet

infotech.link/DF-CheatSheet



Infotech

Infotech
CHEAT SHEET

Infotech OS Data Fabric

Key Concepts & Definitions

Data Lake: Infor's big data management platform.

Data Lake: Flexible and economical cloud object storage solution where data is stored in its raw format.

Data object: Data Lake data is stored as data objects. Data objects are formed from the raw user data and the data object properties.

Atlas: Data Object Explorer (DSE) experience for viewing and managing data lake objects.

Ingestion: A suite of tools that provide data consumers with interfaces for connecting to and processing ANSI SQL queries against data objects stored within the Data Fabric. Supported object formats: REDSHIFT, ORC, CSV, TSV, Par, or user-defined.

Data Lake flow: A sequence of activities orchestrated by CDW that results in landing and/or retaining Data Lake data.

Data cataloging: The Data Catalog stores metadata about data objects that are used within the organization.

Data loader: An application that facilitates the one-time loading of multiple database tables or views already in Data Lake.

Metadata provider: A report that facilitates generating object metadata for tables, views, and materialized views stored in a database.

Data region: The subdomain (partition) of data that traverses Infor's cloud boundary or resides of a job, client, application or system.

Components

Capabilities: The Data Fabric application allows you to view, manage, and query your data objects that are stored in Data Lake. The Data Lake menu includes three menu options: Atlas, Console, and Purge.

API calls:

- RESTful APIs for:
 - Ingesting data into Data Lake
 - Retrieving data from Data Lake or querying Data Lake
 - Managing the data from application
- Message ingestion: Data Lake Ingestion by object, Data Lake Ingestion, Data Lake Storage (see also Ingest Helper)
- Complex JDBC driver: The Complex JDBC driver can be used to query Data Lake data through a third SQL query tool.

Synthetic functions

dataobjectid: Unique identifier of the Data Lake data object.

dataobjectname: The data object name in the source location or file in the info destination.

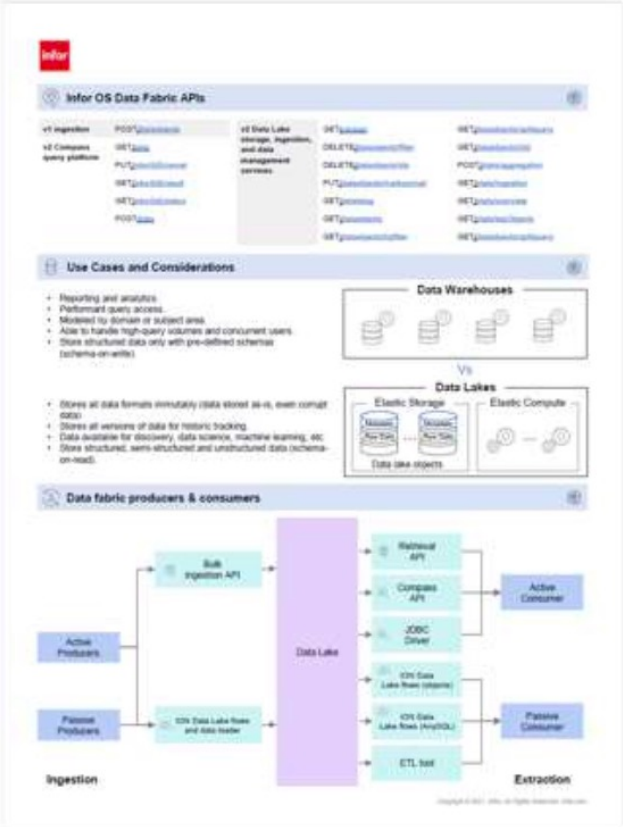
dataobjecttype: Timestamp in which the data object was added to the Data Lake.

dataobjectversion: Detail of variations, including base variation and variations marked as deleted.

dataobjectversionid: Select the maximum variation, even if the highest variation is deleted.

Resources

Product overview, YouTube Playlist, Documentation, Training, Infor technology blog, User Community



Resources

Product Overview, YouTube Playlist, Documentation, Training, Technology Blog, User Community

M3 og Data Fabric

Sanntidstilgang til M3-data gjennom streaming og Stream Pipelines i Data Fabric

Kunstig intelligens og analyse på toppen av Data Lake som **Big Data**-plattform

Lakehouse er Infors **datavarehus** som en tjeneste



Takk

Tid for spørsmål