


INSPARI

presents



Microsoft Fabric

Highlights from Ignite



Speakers:

- **Lars Andersen**, Program Manager, Microsoft
- **Thomas Andersen**, Technology Advisor and Director, Inspari

INSPARI

a valantic company

We deliver end-to-end data, AI & analytics offerings

Strong existing client base - 90% re-occurring revenue

Founded in 2007 with strong growth journey

Top 3 leading Microsoft partner

Technologies



We live our promise: Data Value Guaranteed

7 dedicated full-circle teams

More than 800 customers

30M € in revenue

Locations



Mission

**We empower
decisions,
actions and
change
through people
and data**



INSPARI



Microsoft Fabric is now GA



Status explained

PuPr

Public Preview

Status: test phase

Support: limited

SLA: none

GA

General Availability

Status: final

Support: normal

SLA: yes



Microsoft Fabric is ready for production




What is Microsoft Fabric

Data analytics for the era of AI

Microsoft Fabric is an **all-in-one analytics solution** for enterprises that covers everything from data movement to data science, Real-Time Analytics, and business intelligence. It offers a comprehensive suite of services, including **data lake, data engineering, and data integration, all in one place.**

With Fabric, you don't need to piece together different services from multiple vendors. Instead, you can enjoy a highly integrated, end-to-end, and easy-to-use product that is **designed to simplify your analytics needs.**



Key concepts

**Complete
Analytics
Platform**

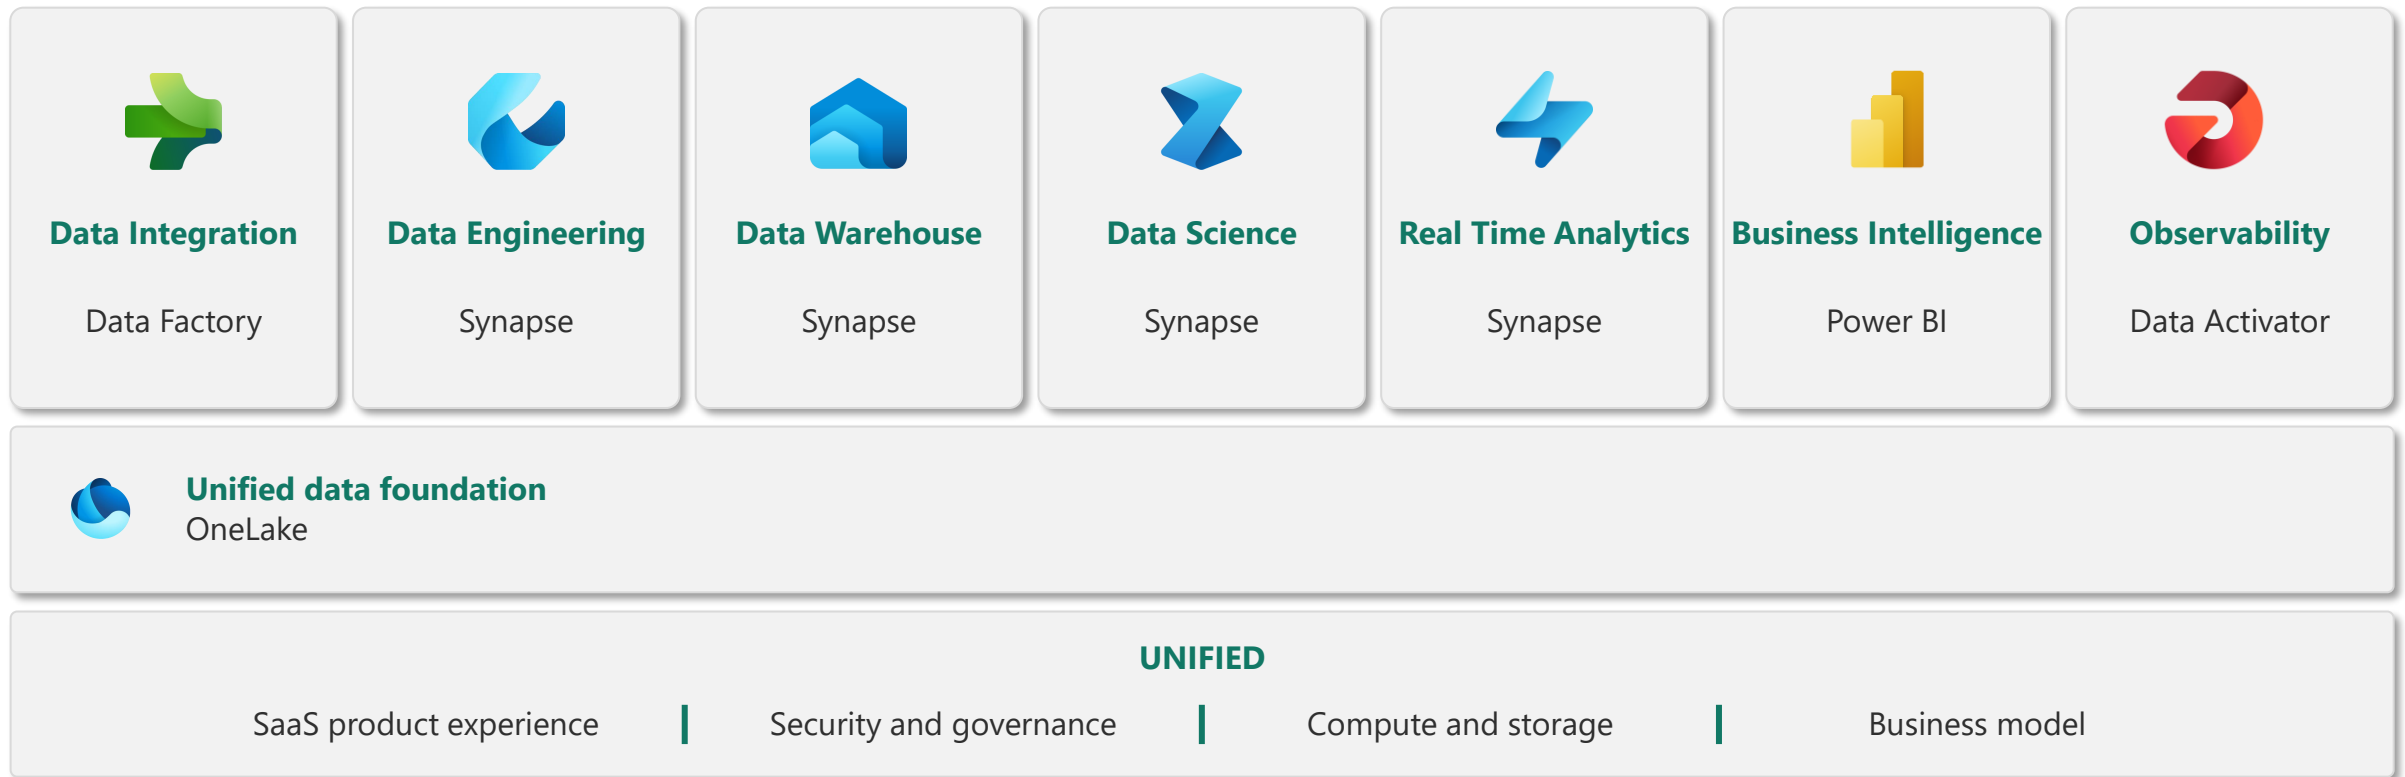
**Lake Centric
and Open**

**Empower Every
Business User**

**AI
Powered**



Microsoft Fabric at a glance



Software as a Service

Easy and accessible

Frictionless onboarding

Instant provisioning

Quick results w/ intuitive UX

Success by default

Minimal knobs

Auto optimized

Auto integrated

Centralized administration

Tenant-wide governance

Centralized security

Compliance built-in






Data Engineers

- **Execute faster** with the ability to spin up a Spark VM cluster in seconds, or configure with familiar experiences like Git DevOps pipelines for data engineering artifacts
- **Streamline your work** with a single platform to build and operate real-time analytics pipelines, data lakes, lake houses, warehouses, marts, and cubes using your preferred IDE, plug-ins, and tools.
- **Reduce costly data replication** and movement with the ability to produce base datasets that can serve data analysts and data scientists without needing to build pipelines

Serve data via warehouse or lakehouse

Supporting experiences:




Data Scientists

- **Quickly tune a custom model** by integrating a model built and trained in Azure ML in a Spark notebook
- **Work faster** with the ability to use your preferred data science frameworks, languages, and tools
- **Bypass engineering dependencies** with the ability to use your preferred no-code ML Ops to deploy and operate models in production
- **Tap into proven-at-scale models and services** to accelerate your AI differentiation (AOAI, Cognitive Services, ONNX integration, etc).

Serve transformed data

Supporting experiences



Data Analysts

- **Avoid slow, progress-stagnating data wrangling** by seamlessly triggering a workflow that can unlock data engineering tools and capabilities quickly.
- **Accelerate your work** with visual and SQL based tools for self-serve data transformations and modeling as well as self-serve tools for reporting, dashboards, and data visualizations
- **Turn data into impact** with industry-leading BI tools and integration with the apps your people use everyday like Microsoft 365

Serve insights via embedding

Supporting experiences



Data Citizens

- **Make more data-driven decisions** with actionable insights and intelligence in your preferred applications
- **Maintain access to all the data you need**, without being overwhelmed by data ancillary to your role thanks to fine grain data access management controls

Supporting experiences



Data Stewards


- **Maintain visibility and control of costs** with a unified consumption and cost model that provides evergreen spend optics on your end-to-end data estate
- **Gain full visibility and governance** over your entire analytics estate from data sources and connections to your data lake, to users and their insights

Serve data via warehouse or lakehouse

Data Engineers



Supporting experiences:



Data Scientists



Supporting experiences



Data Analysts




Supporting experiences



Data Citizens



Supporting experiences



INSPIRI

Serve data via warehouse or lakehouse

Serve transformed data

Serve insights via embedding

Serve data via warehouse or lakehouse

OneLake




One data lake for the entire organization
One copy of data for use with multiple analytical engines



Data Stewards

- **Maintain visibility and control of costs** with a unified consumption and cost model that provides evergreen spend optics on your end-to-end data estate
- **Gain full visibility and governance** over your entire analytics estate from data sources and connections to your data lake, to users and their insights

Shortcuts





Highlights

Public Preview

Copilot for Power BI, Data Factory, Data Engineering, Data Science

General Available

Git + Deployment Pipelines for many artifacts (notebook, lakehouse, warehouse)

Public Preview

REST APIs for many artifacts

Coming Soon

Mirroring for Warehouse

Public Preview


Support for vnet gateway in Dataflow Gen2

General Available

Fabric Data Factory feature parity with Azure Data factory

General Available

OneLake/Lakehouse connector in Azure Data Factory





Power BI

Public Preview

Copilot for Power BI

General Available

Direct Lake now supports all data in OneLake, including data from Data Warehouse

Public Preview


Stored Credentials for Direct Lake semantic model RLS/OLS security

Public Preview

OneLake integration for import models

Public Preview

Explore





Data Warehouse

General Available

Query Insights – 30-day query log

General Available

Publishing full DML to Delta Lake logs – max 1 min. latency in delta lake

General Available

Automatic Data Compaction behind the scenes

General Available

Improvements to CSV data ingestion – now in Fabric Warehouse

General Available

Read multi-TB results – buffering and streaming large result sets

General Available

Fast compute resource assignment – compute resources in milliseconds

General Available

SSD Metadata Caching – file and rowgroup metadata now cached

General Available

Time Travelling using Table Clones

General Available

SQLPackage – database automation

General Available

REST- APIs access

General Available

Dynamic Data Masking

Coming Soon

Warehouse Git-enabled (in the upcoming weeks)



Data Engineering

General Available

REST API support for Lakehouse

Public Preview

REST API support for Spark Job Definition

Public Preview

REST API support for Notebooks

Public Preview

Lakehouse support for git integration and deployment pipelines

Public Preview

Notebook support for git integration

Public Preview

Notebook in Deployment Pipeline

Public Preview

Environment

General Available

Notebook resources .JAR file support

General Available


Mssparkutils new API – reference run multiple notebooks in parallel

General Available

Upgraded DataGrid in Lakehouse – sorting, filtering, resizing columns

General Available

SQL Endpoint re-provisioning in Lakehouse





Data Science

Public Preview

Copilot in Notebooks

General Available

Custom Python Operations in Data Wrangler

Public Preview

Data Wrangler for Spark DataFrames

General Available

MLFlow Notebook Widget

General Available

New Model & Experiment Item Usability Improvements

General Available

Recent Experiment Runs overview

General Available

Models renamed to ML Models

General Available

Release of SynapseML v1.0

General Available

Train Interpretable Explainable Boosting Machines with SynapseML

General Available

Prebuilt AI models

General Available

Reusing Existing Spark Session in sparklyr

General Available

Rest API Support for ML Experiments and ML Models



Data Factory – Dataflow Gen2

General Available

Fabric connectors

General Available

Automatic refresh cancellation

General Available

Error message propagation through gateway

General Available

Support for column binding for SAP HANA connector

General Available


Staging artifacts will be hidden

General Available

New Power Query Editor style

Public Preview

Support for VNET Gateways





Data Factory – Pipelines

General Available

Cross workspace “Save as”

General Available

Dynamic content flyout integration with Email and Teams activity

General Available

Copy activity now supports fault tolerance for Fabric Data Warehouse connector

General Available

MongoDB and MongoDB Atlas connectors are now available

Public Preview

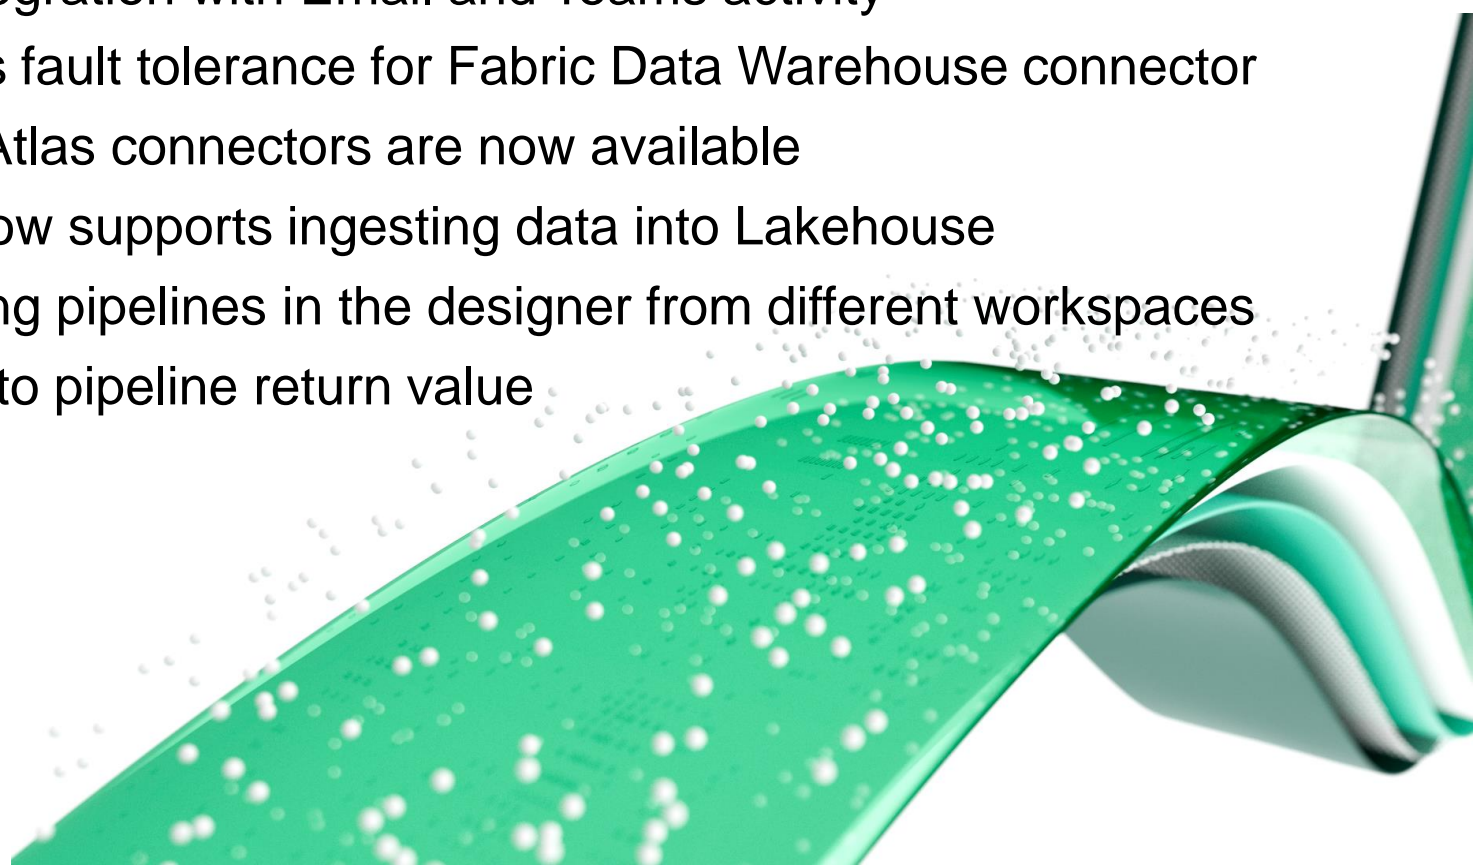
Microsoft 365 connector now supports ingesting data into Lakehouse

General Available

Multi-task support for editing pipelines in the designer from different workspaces

General Available

String interpolation added to pipeline return value



Your next questions

1. How does Fabric fit into our existing solution?
2. Can Fabric coexist with our current platform?
3. Can we transition to Fabric or is it a big bang implementation?
4. ...

Your next action

Contact Inspari for a meeting to talk about concerns and ambitions, next step could be a workshop for further clarification or a pilot where we test it out



Leading people – leading business

INSPARI



Thank you

Contact: tha@inspari.dk