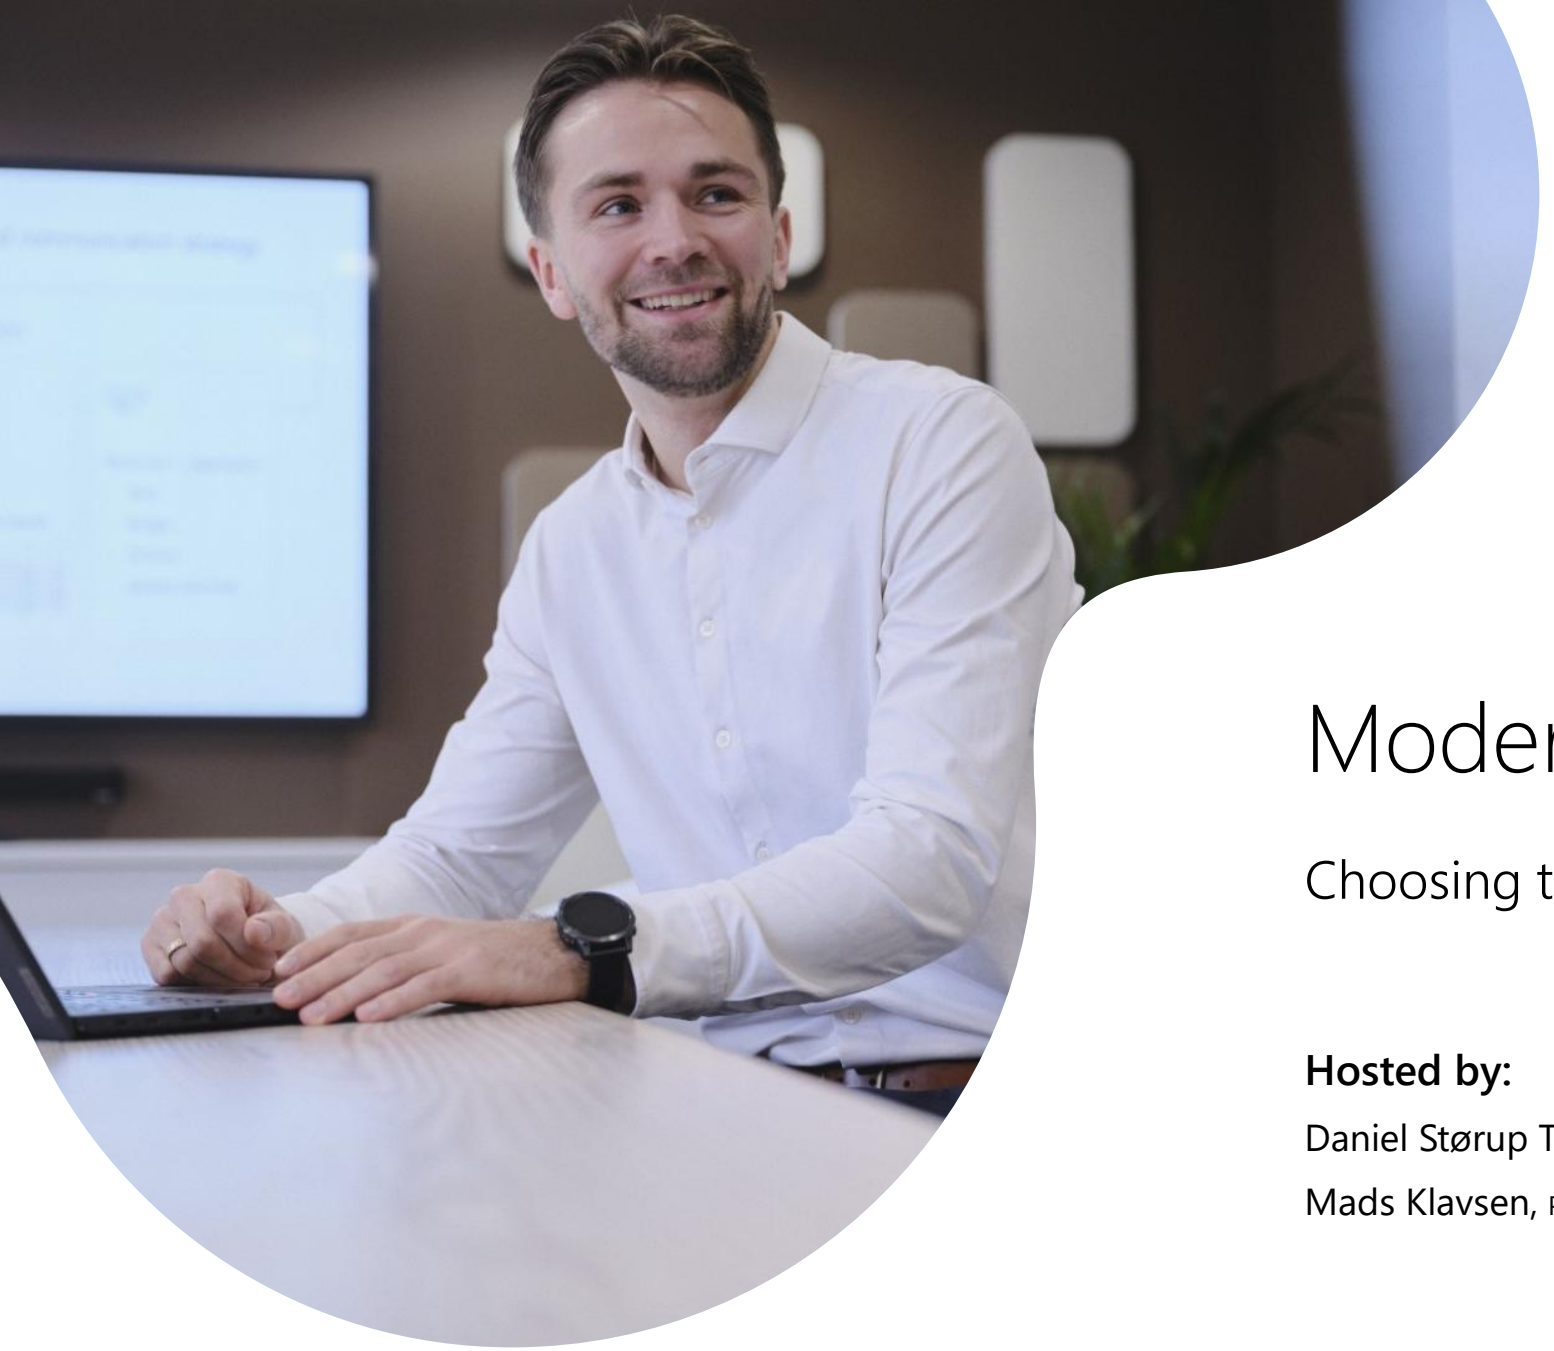

Modern Data Platform

Choosing the right fit for an AI-ready future

Hosted by:

Daniel Størup Thomsen, Senior Data Platform Architect

Mads Klavsen, Principal Data Engineer

Today's Presenters

DANIEL STØRUP THOMSEN

Senior Data Platform Architect
Capability Expert, Databricks

MADS KLAVSEN

Principal Data Engineer
Capability Expert, Fabric

Practical Information

The webinar is being recorded. Recording and slides will be shared after the session.

Q&A will be held at the end of the presentation. Please write your questions in the chat.

INSPARI AT A GLANCE

Your partner for business impact with AI & data

Our mission

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

Strong team

200+ data & AI experts

Copenhagen, Aarhus & Bucharest

Comprehensive capabilities for successful AI & data transformation

- Data & AI Strategy
- Data & AI Architecture
- Data & AI Products
- Data & AI Management & Governance
- Data & AI Engineering
- Data & AI Ops

Data Value Guaranteed

Recognized thought leader

Podcast on data, AI, and digital transformation

Yearly conference with 500 attendees

Webinars & events

Microsoft MVP within data platform

A well-built Microsoft proficiency

Fabric Featured Partner

Cloud Solution Provider

Azure

Microsoft Solutions Partner
Digital & App Innovation
Azure

Microsoft Solutions Partner
Business Applications

Microsoft Solutions Partner
Infrastructure
Azure

Microsoft Solutions Partner
Data & AI
Azure

Specialist
Analytics
AI and Machine Learning

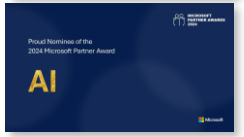

Substantial track record of successful AI & data implementations

PANDORA · ENERGINET · Dansk Industri · BESTSELLER · SIEMENS · CPH · midt regionmidtjylland · bankdata · nne novo nordisk engineering

Selection of technology partners

snowflake
Growth Partner of the Year 2025

databricks

INSPARI
a valantic company

VALANTIC FACTS & FIGURES

We are part of valantic – a European leader within digital transformation

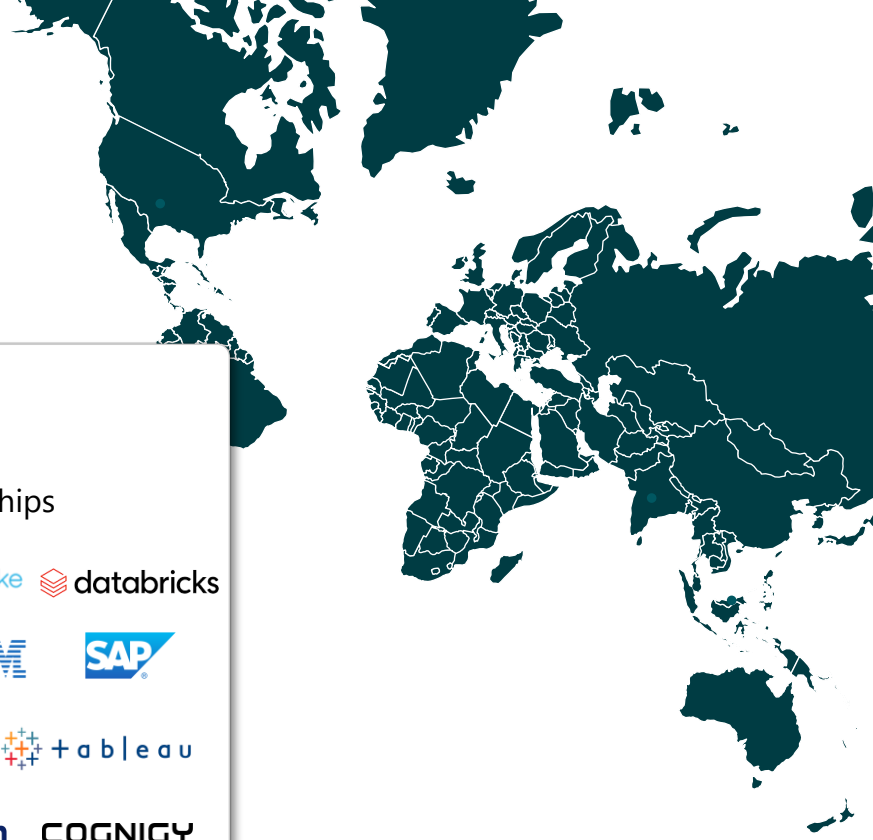

600+
experts in data & AI globally

4000+
digital experts globally

valantic offers digitalization know-how for every customer challenge and provide comprehensive, end-to-end support

+500
Blue Chip Clients

Automotive & mobility

Porsche, BMW, Toyota, MAERSK, Lufthansa

Pharma, medtech, chemicals

novo nordisk, Pfizer, BAYER, AstraZeneca

Retail & fashion

SWAROVSKI, PANDORA, Red Bull, BESTSELLER

Manufacturing

BOSCH, BioMar, SIEMENS, ABB

Telco & utilities

ENERGINET, Ørsted, vodafone

Banking & insurance

Deutsche Bank, AOK, Santander

+20
Strategic Tech Partnerships

Microsoft, snowflake, databricks, aws, HIVEMQ, IBM, SAP, Informatica, UiPath, tableau, CCH Tagetik, Anaplan, COGNIGY, SIGNAVIO, data iku, celonis, outsystems, shopify, Spryker, ivalua, Wolters Kluwer, Adobe, CAMUNDA, kinaxis, commercetools, salesforce, PIMCORE, shopware

INSPARI
a valantic company

Agenda

- What is a modern data platform
- Platform architecture and overview
- Comparison of platforms
- AI-ready platforms and best of breed

What Do We Mean by A Modern Data Platform?

A successful data platform requires alignment between strategy, people, processes, and technology.

Strategy:
Why are we doing this
& how will we succeed?

Operating Model:
How are we going to
execute & with what?

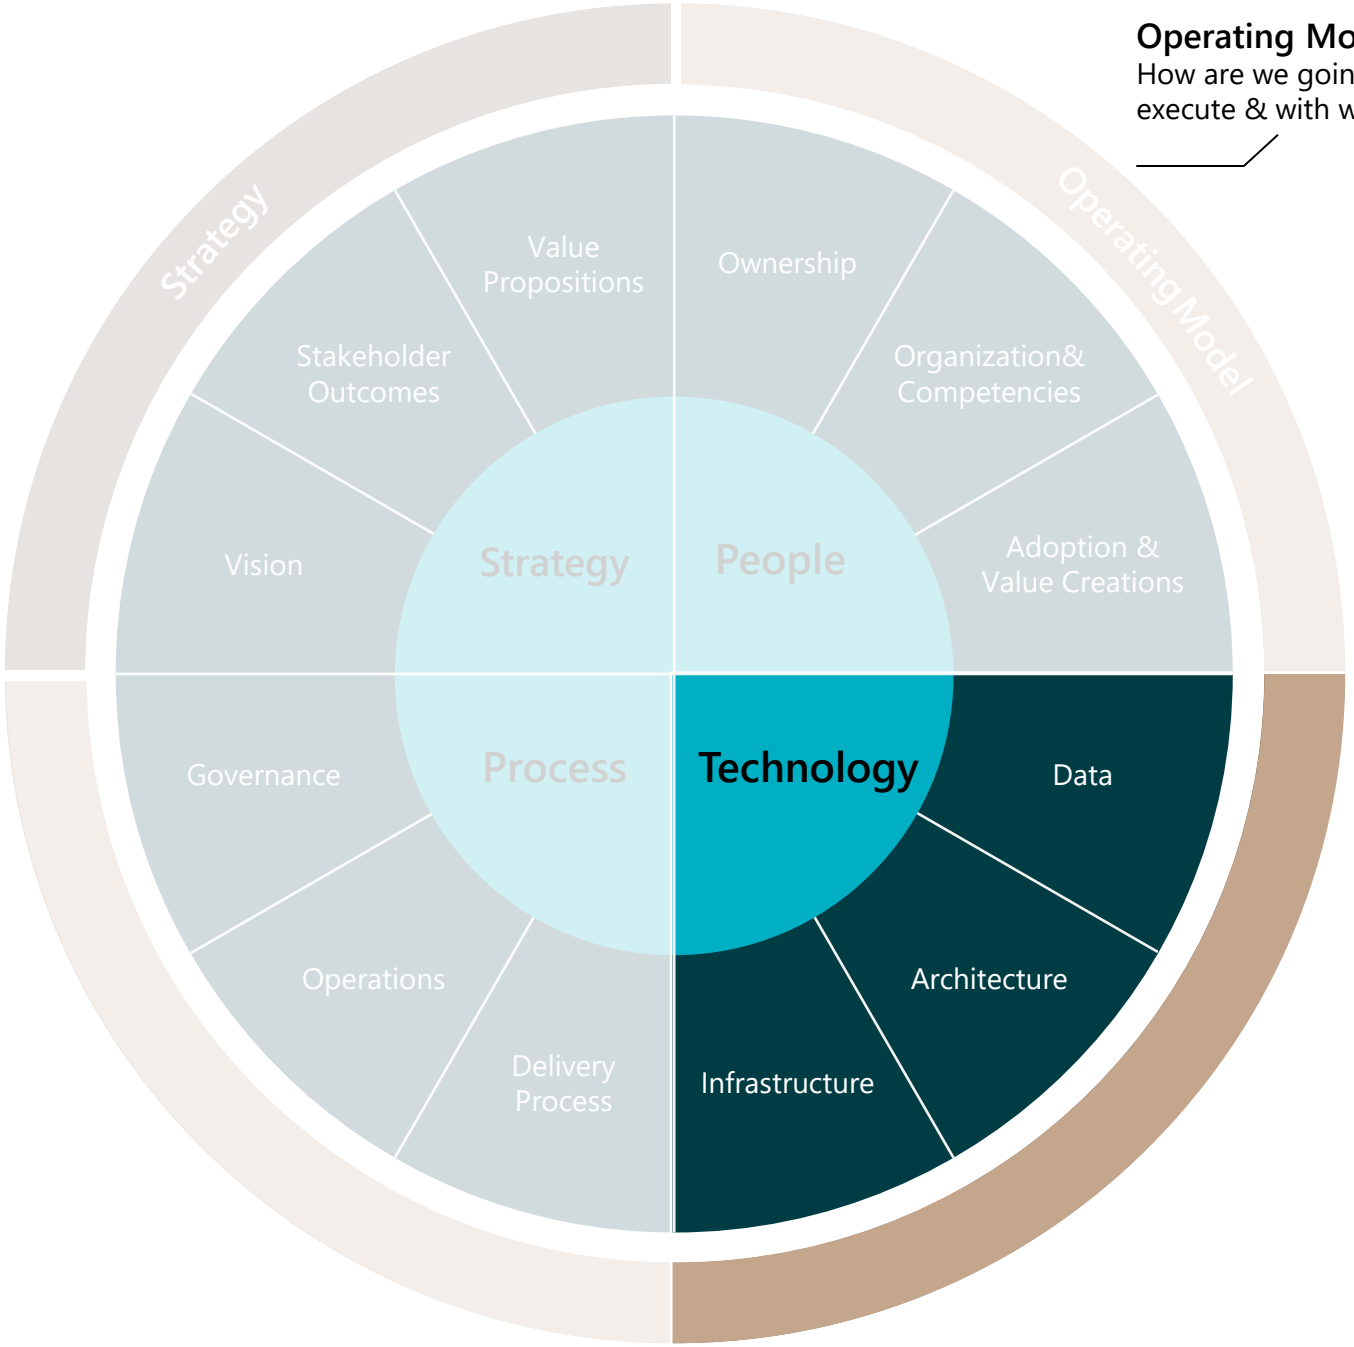

The Data Platform Landscape is Complex

Open Formats & Ecosystems

Open Table Formats (delta, Iceberg)

Ecosystem integrations

Interoperability between tools

Overlapping Capabilities

Data Engineering

Analytics

AI/ML

Governance

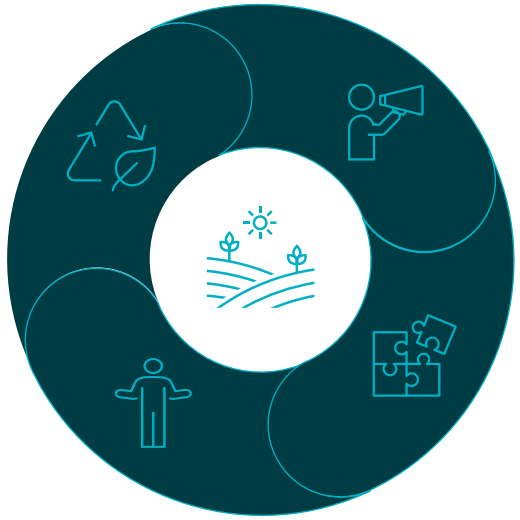

Popular Platform Terminology

Unified platform

AI-ready

Lakehouse

Rapid Innovation

New features constantly

Changing names and architectures

Customization or platform-native

Focusing on Four Platforms

Microsoft Fabric

Snowflake

Databricks

Best of Breed

From Data Warehouses to Unified Analytics Platform

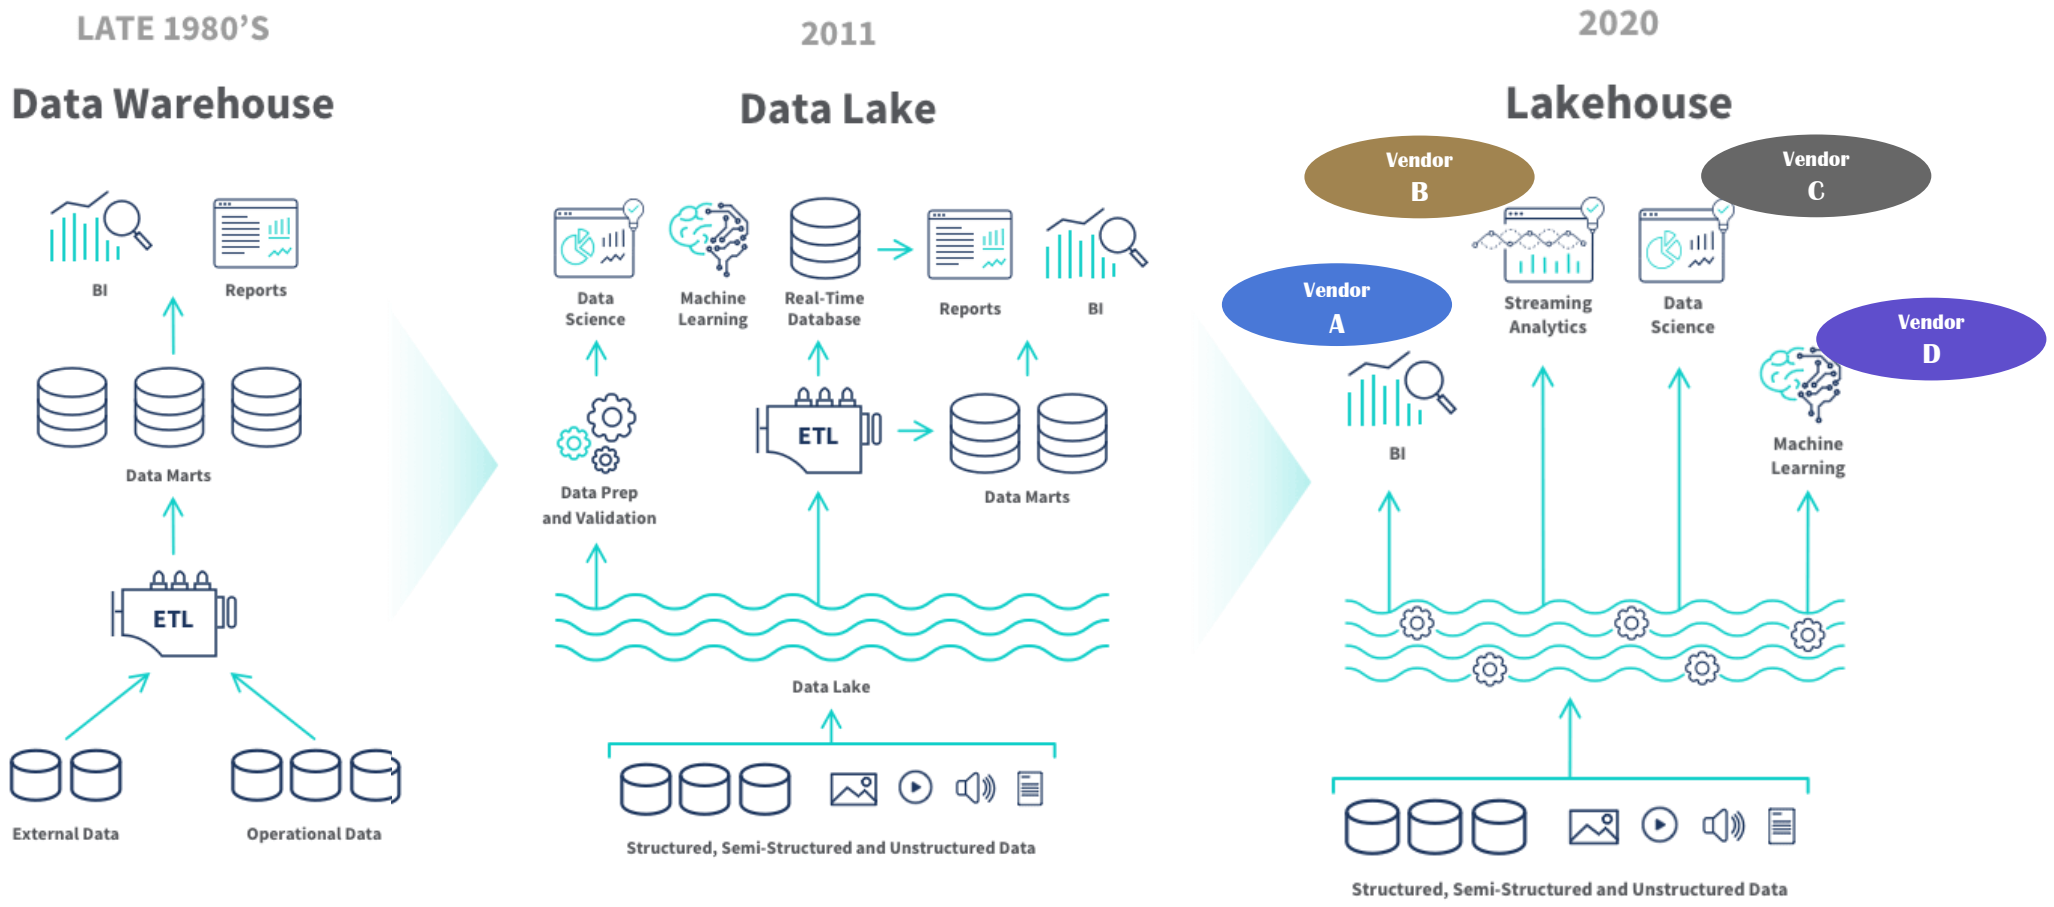

From Data Warehouses to Unified Analytics Platform

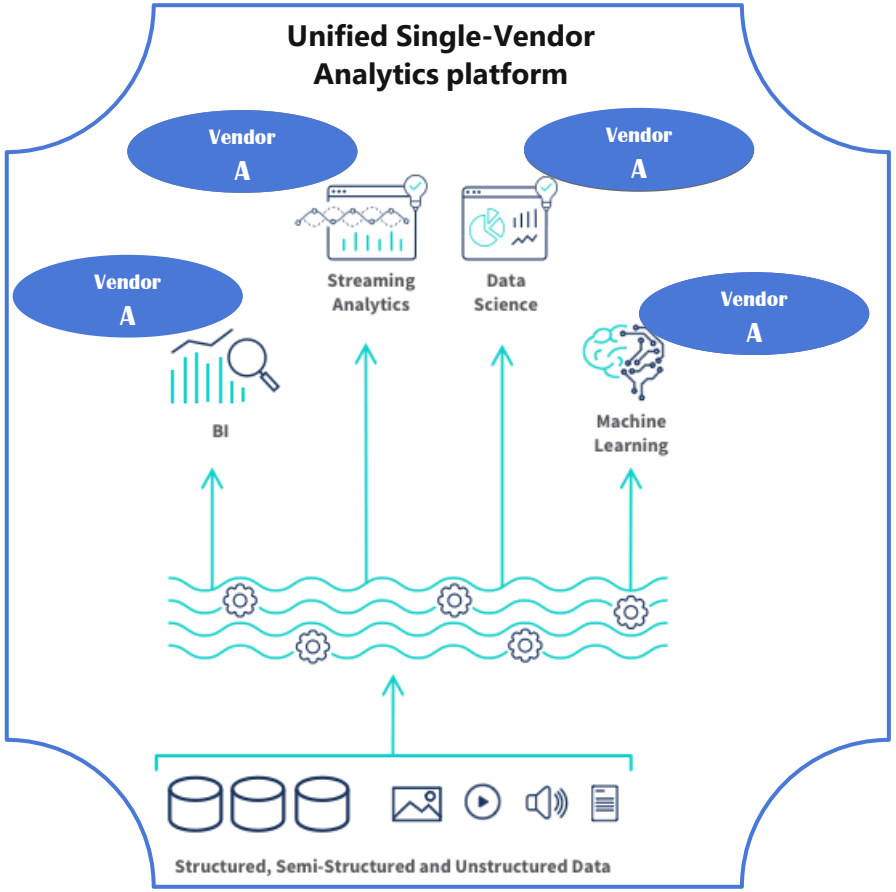

Solving the Same Problem

Storage

Data Lake
Data Warehouse
Lakehouse
Open Table Formats

Data Ingestion

Applications
Databases
APIs
Streams

Processing

Data Pipelines
Transformations
Data Engineering

Data Quality

Accuracy
Consistency
Monitoring

Analytics / AI

Dashboards
Self-service
Machine Learning
AI capabilities

Governance

Catalog
Lineage
Security

Agenda

- What is a modern data platform
- **Platform architecture and overview**
- Comparison of platforms
- AI-ready platforms and best of breed

Data ingestion & processing (batch, streaming, SDP)

Lakehouse architecture (delta lake/Iceberg)

Notebooks + Jobs (development + orchestration)

Delta sharing & AI/BI

Unity Catalog - Governance

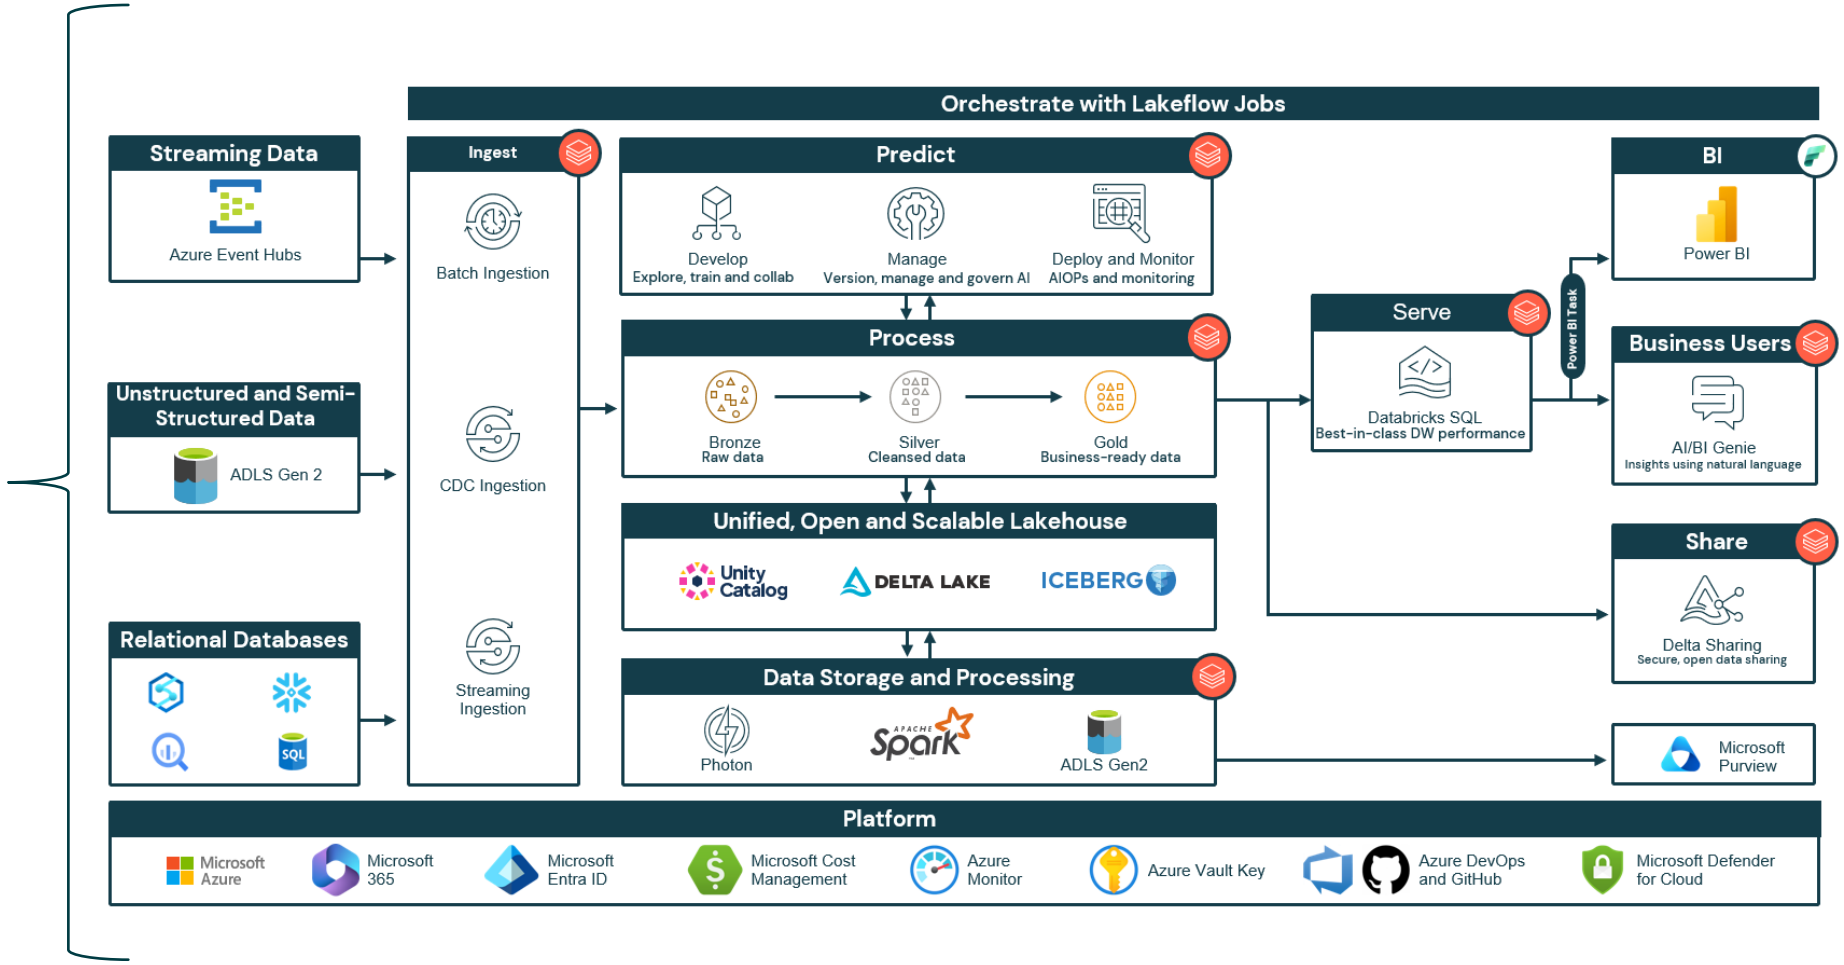

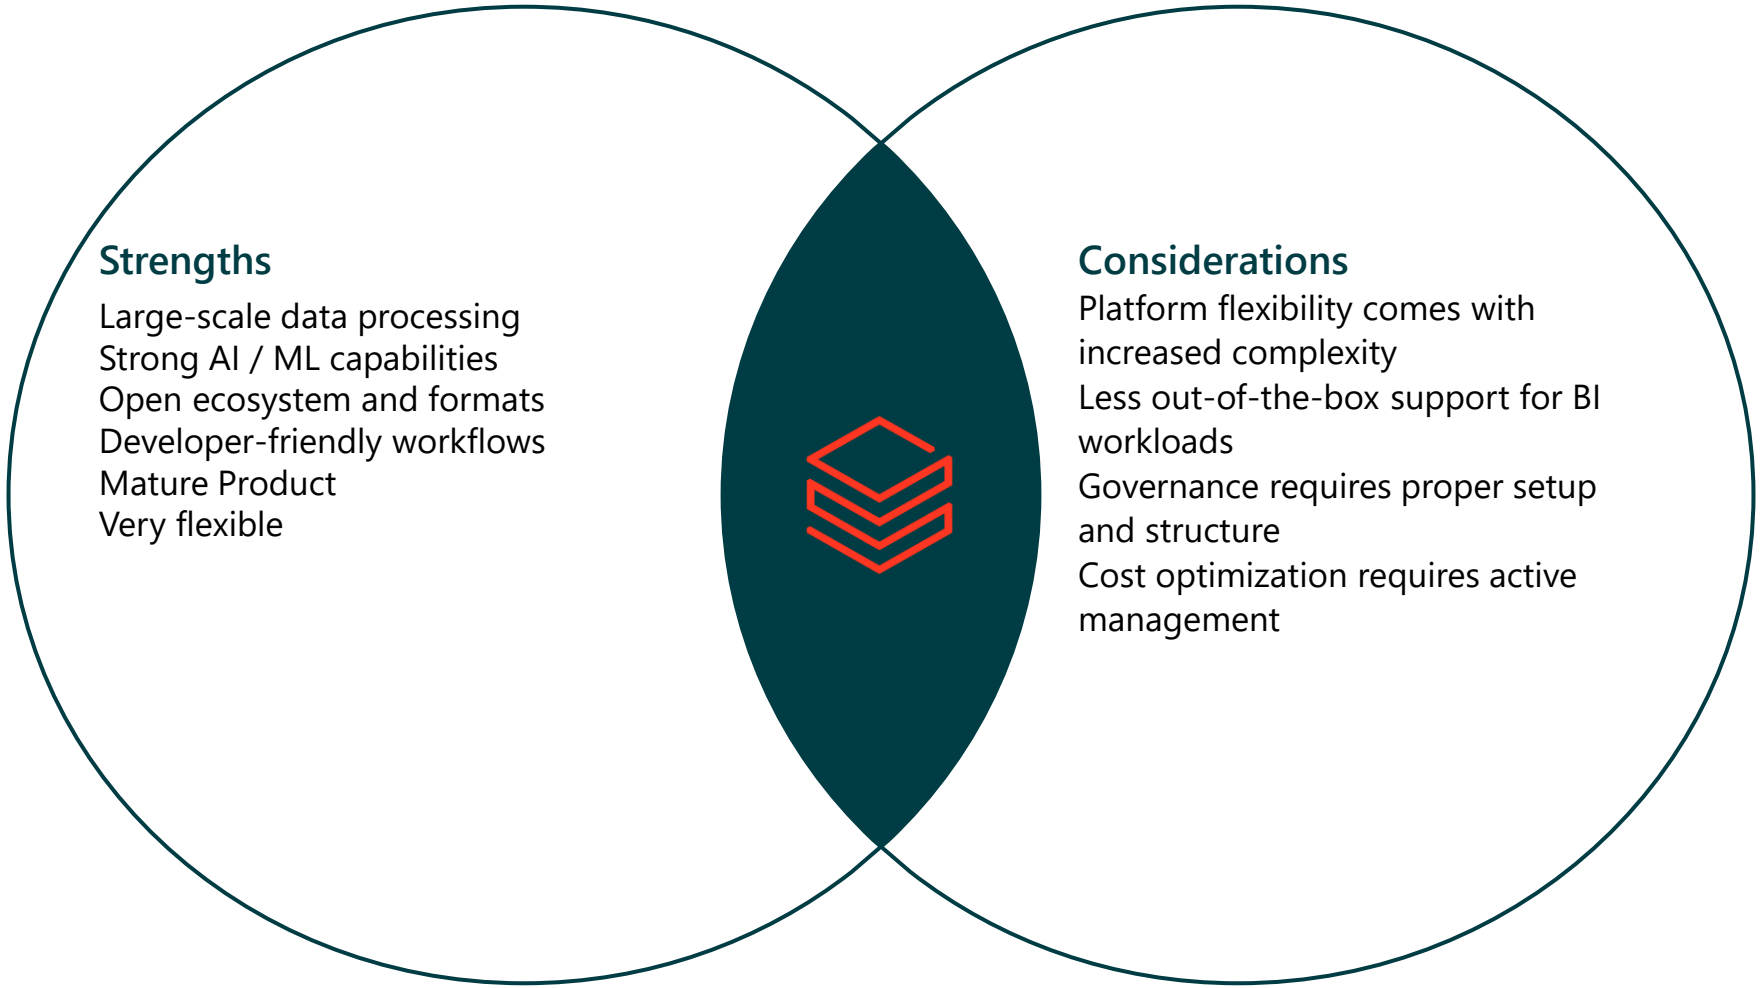
Strengths
Large-scale data processing
Strong AI / ML capabilities
Open ecosystem and formats
Developer-friendly workflows
Mature Product
Very flexible
Considerations
Platform flexibility comes with increased complexity
Less out-of-the-box support for BI workloads
Governance requires proper setup and structure
Cost optimization requires active management

Cross-cloud foundation

Separation of storage and compute

Multi-language compute layer

Built-in governance at the core (Horizon Catalog)

Strong data sharing and ecosystem integration

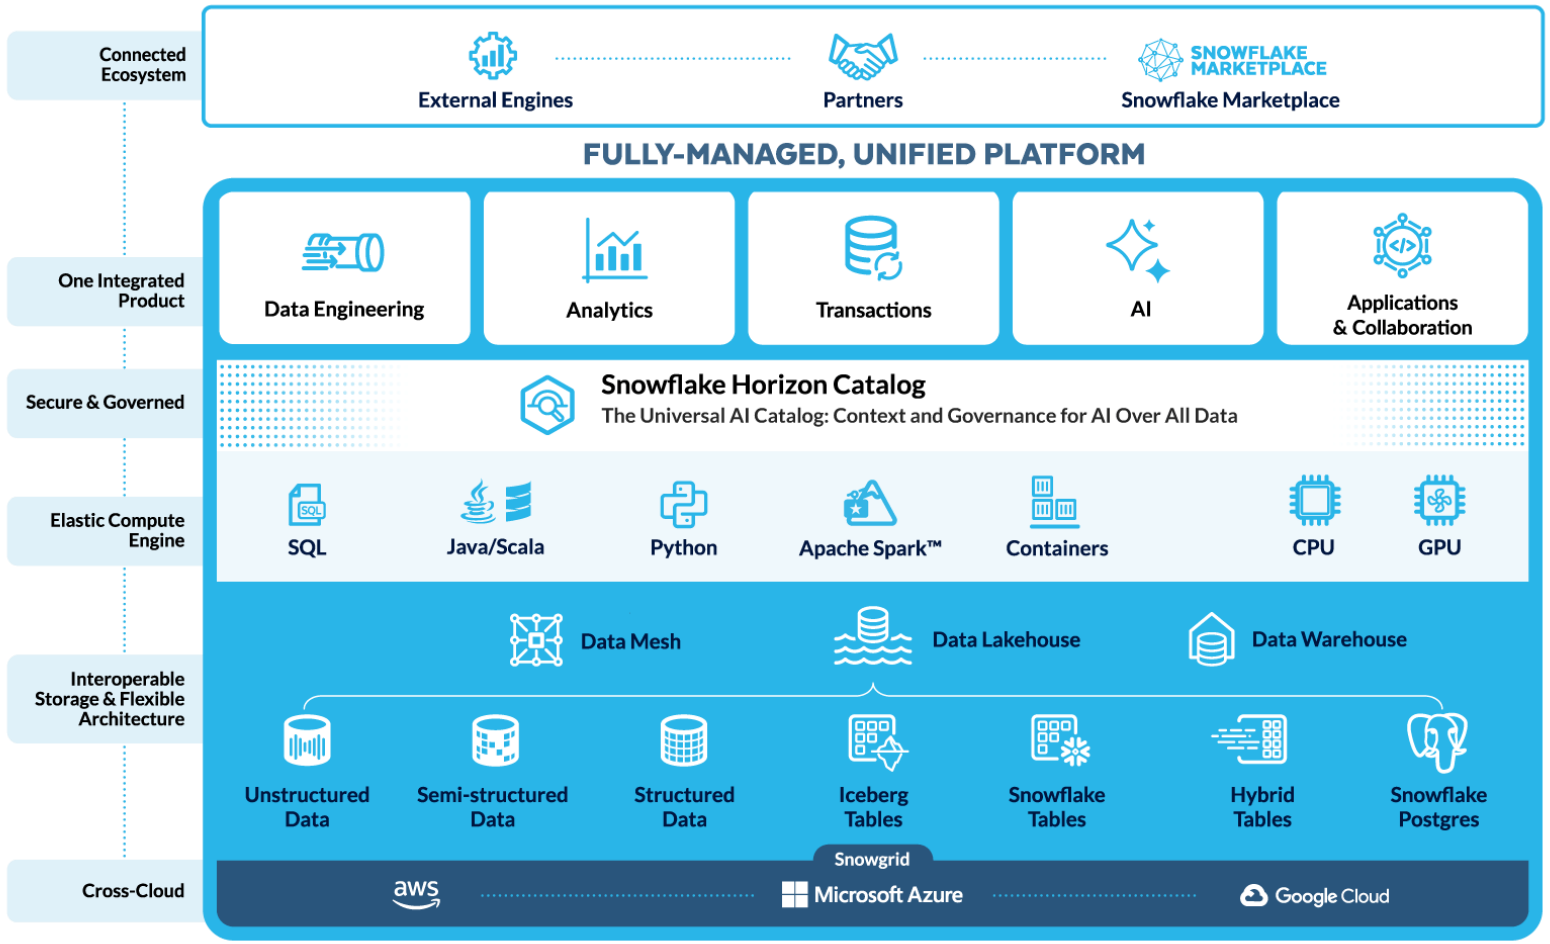

INSPARI
a valantic company

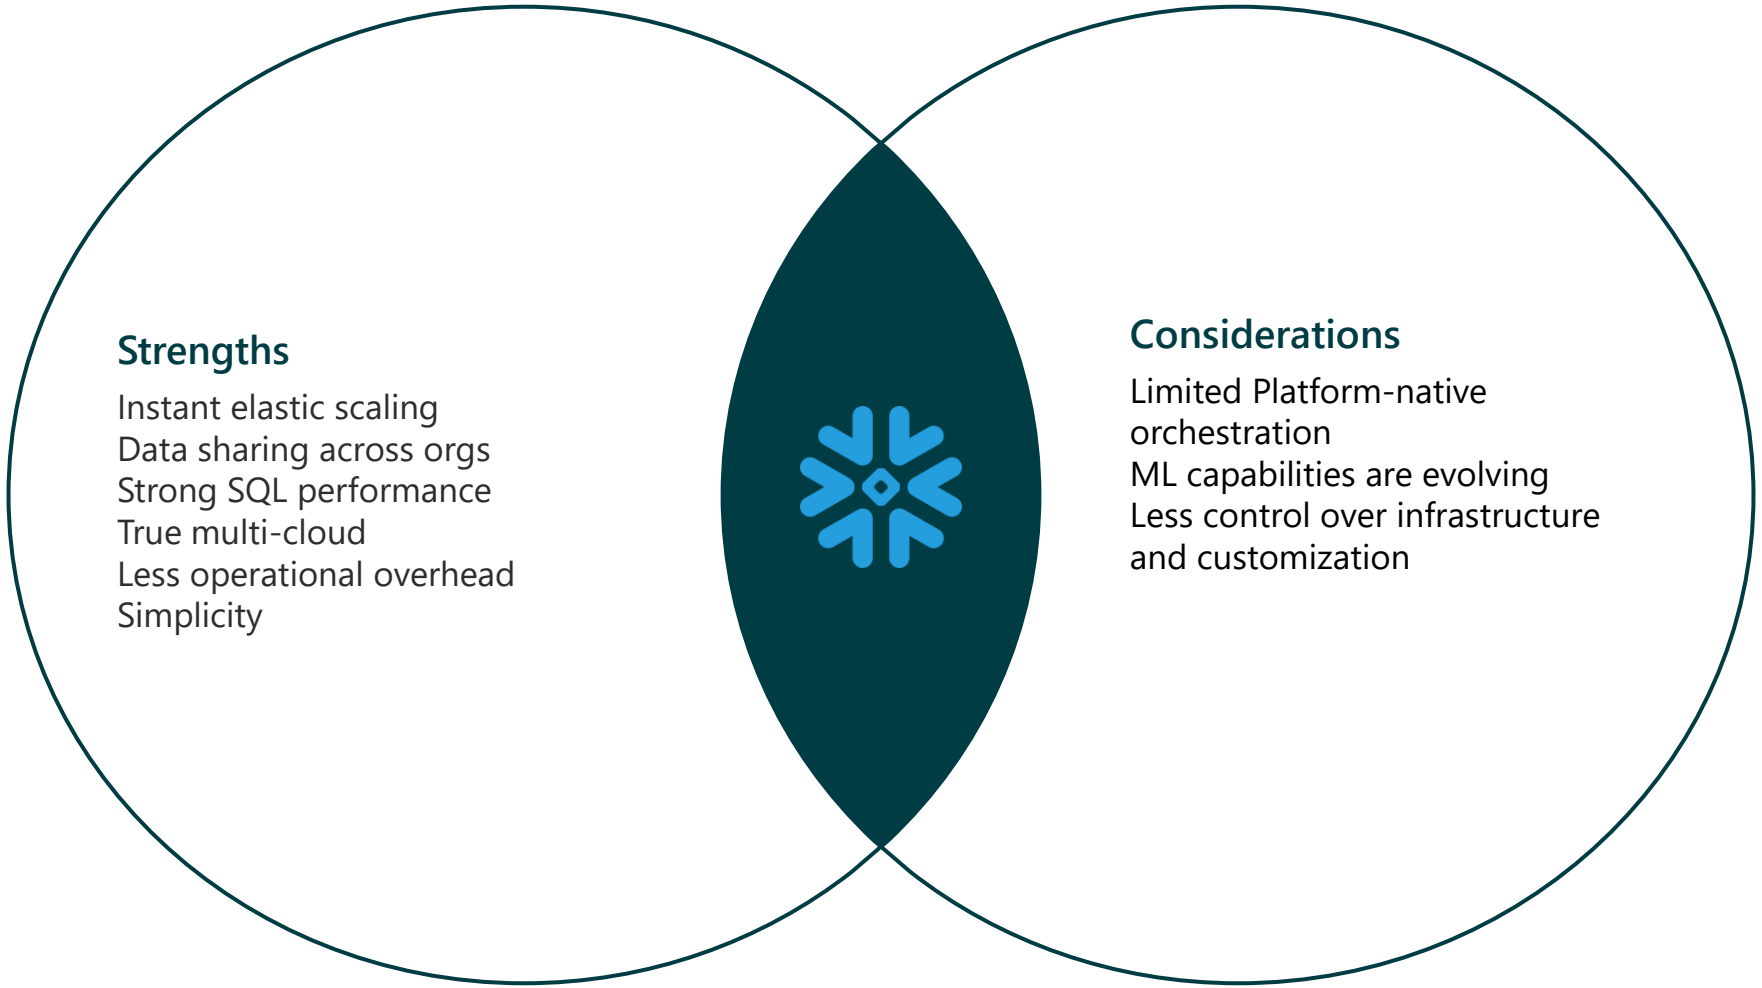
Strengths
Instant elastic scaling
Data sharing across orgs
Strong SQL performance
True multi-cloud
Less operational overhead
Simplicity
Considerations
Limited Platform-native orchestration
ML capabilities are evolving
Less control over infrastructure and customization

Different workloads – same platform

OneLake – One unified storage

Lakehouse + Data Warehouse

No-code vs notebooks

Power BI – Direct Lake

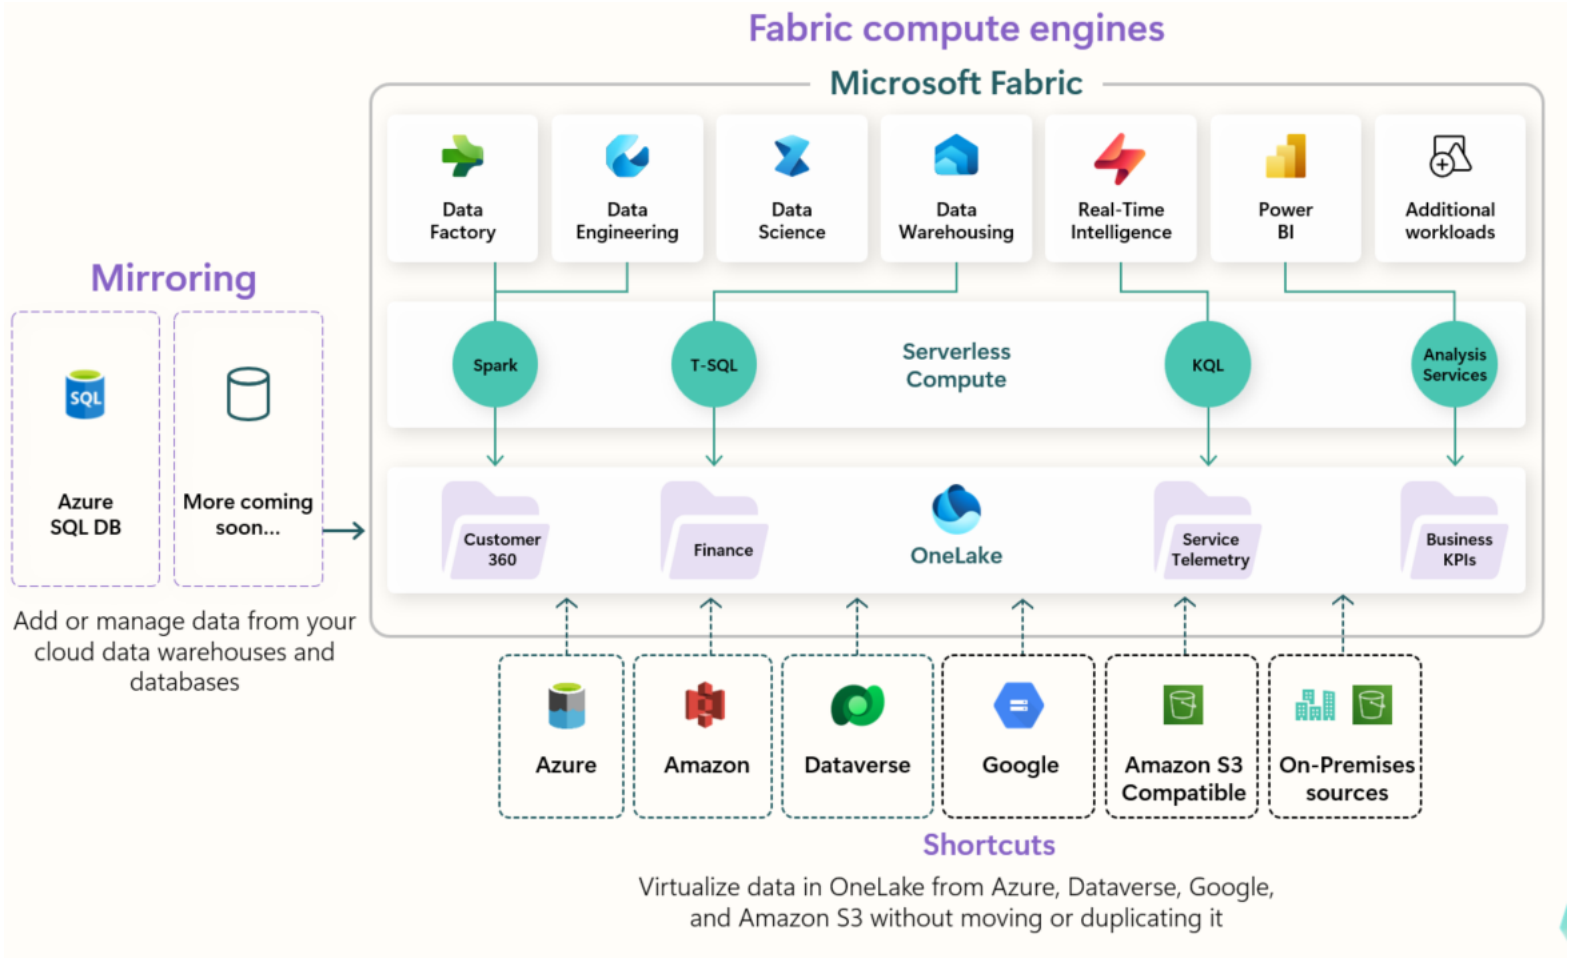

INSPARI
a valantic company

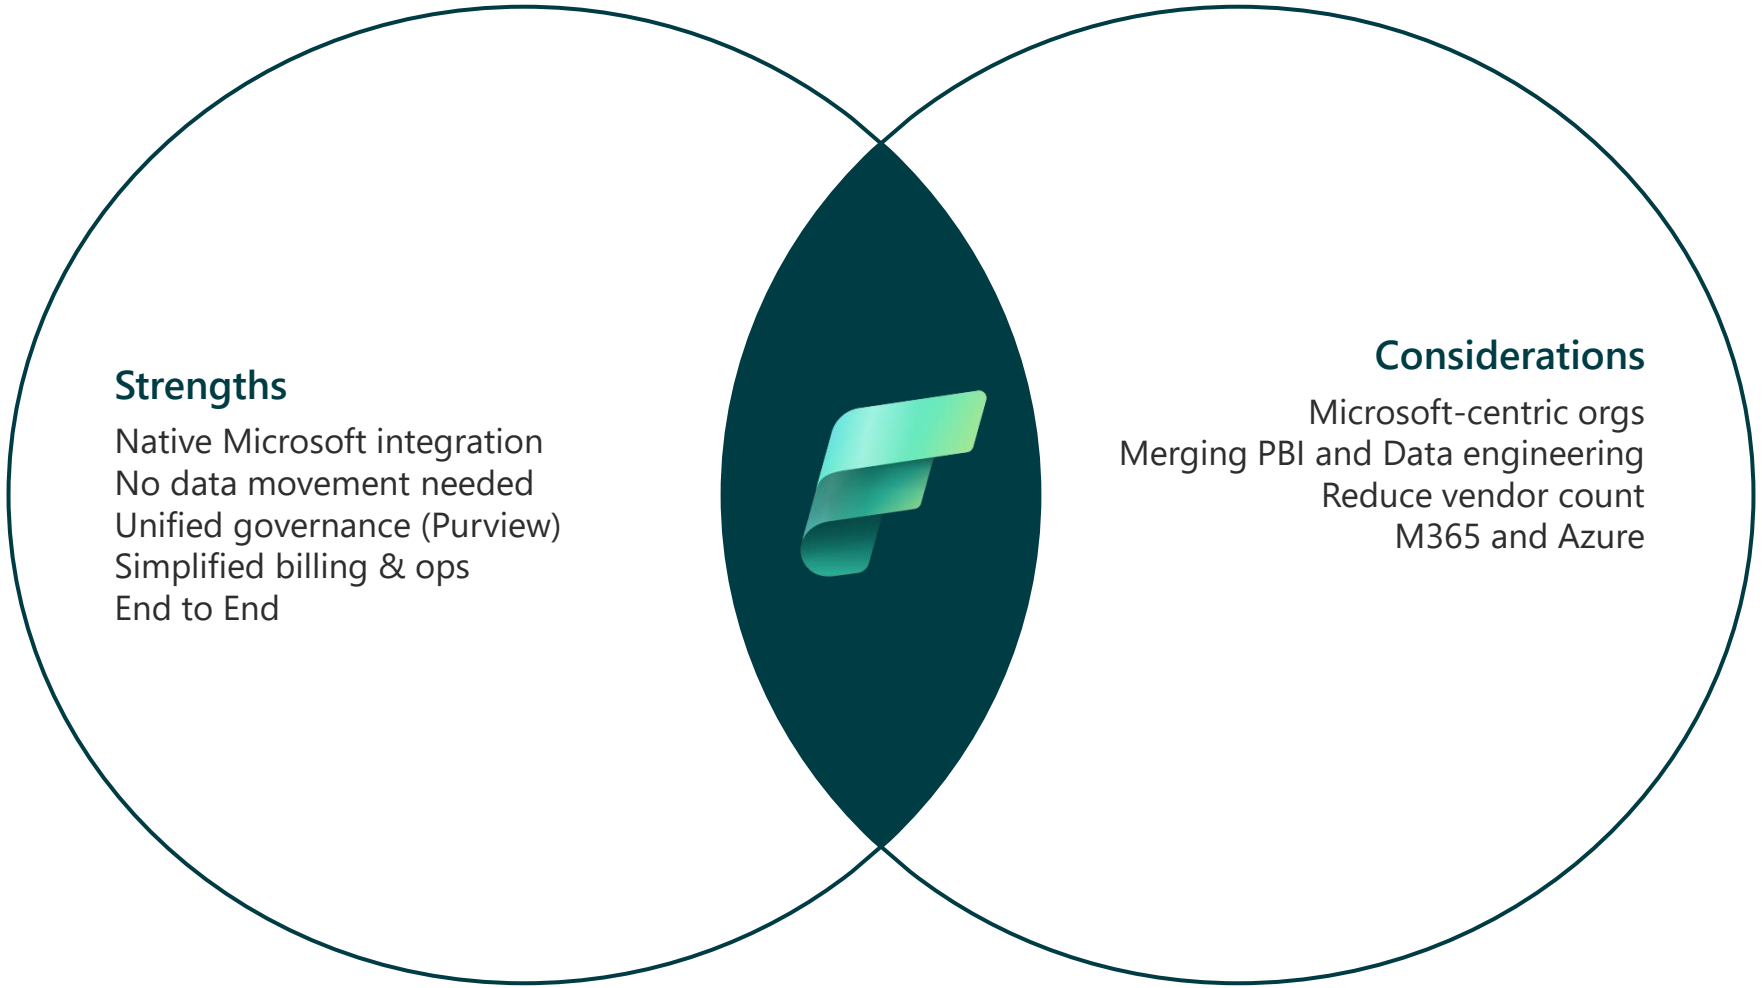
Strengths
Native Microsoft integration
No data movement needed
Unified governance (Purview)
Simplified billing & ops
End to End
Considerations
Microsoft-centric orgs
Merging PBI and Data engineering
Reduce vendor count
M365 and Azure

Agenda

- What is a modern data platform
- Platform architecture and overview
- **Comparison of platforms**
- AI-ready platforms and best of breed

Value and success criteria definition

Selection of appropriate use cases for testing

Project infrastructure setup to test selected platform solutions

Develop, execute, and benchmark platform solutions

Synthesize and recommend future data platform

Solution Assessment

Foundation for decision when selecting vendor and technology to enable your vision.

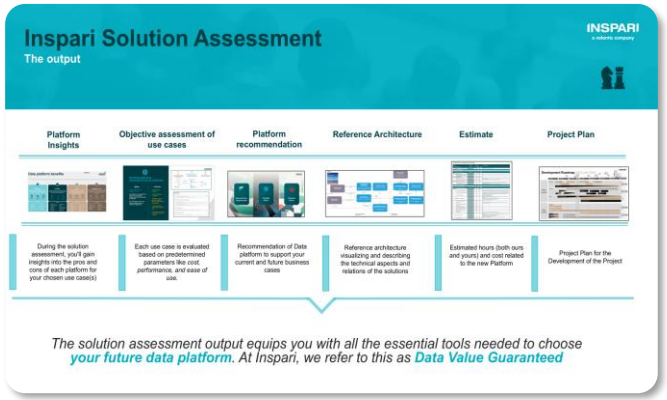

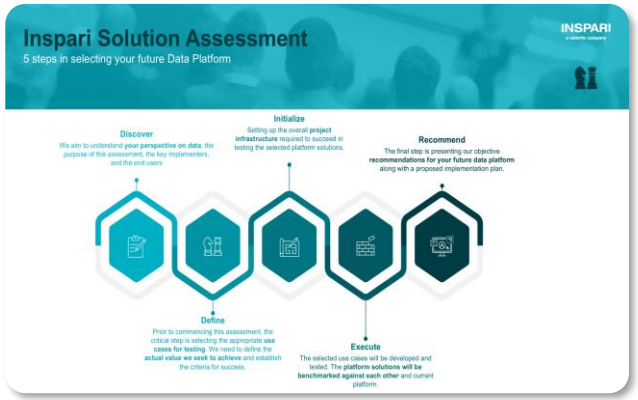

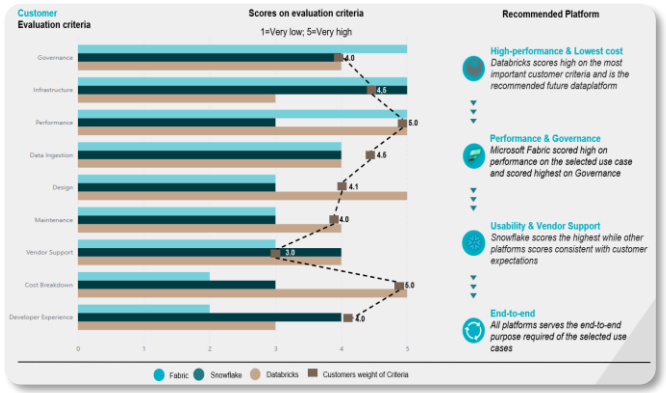

Evaluation Criteria

Evaluation criteria

Scores on evaluation criteria

1=Very low; 5=Very high

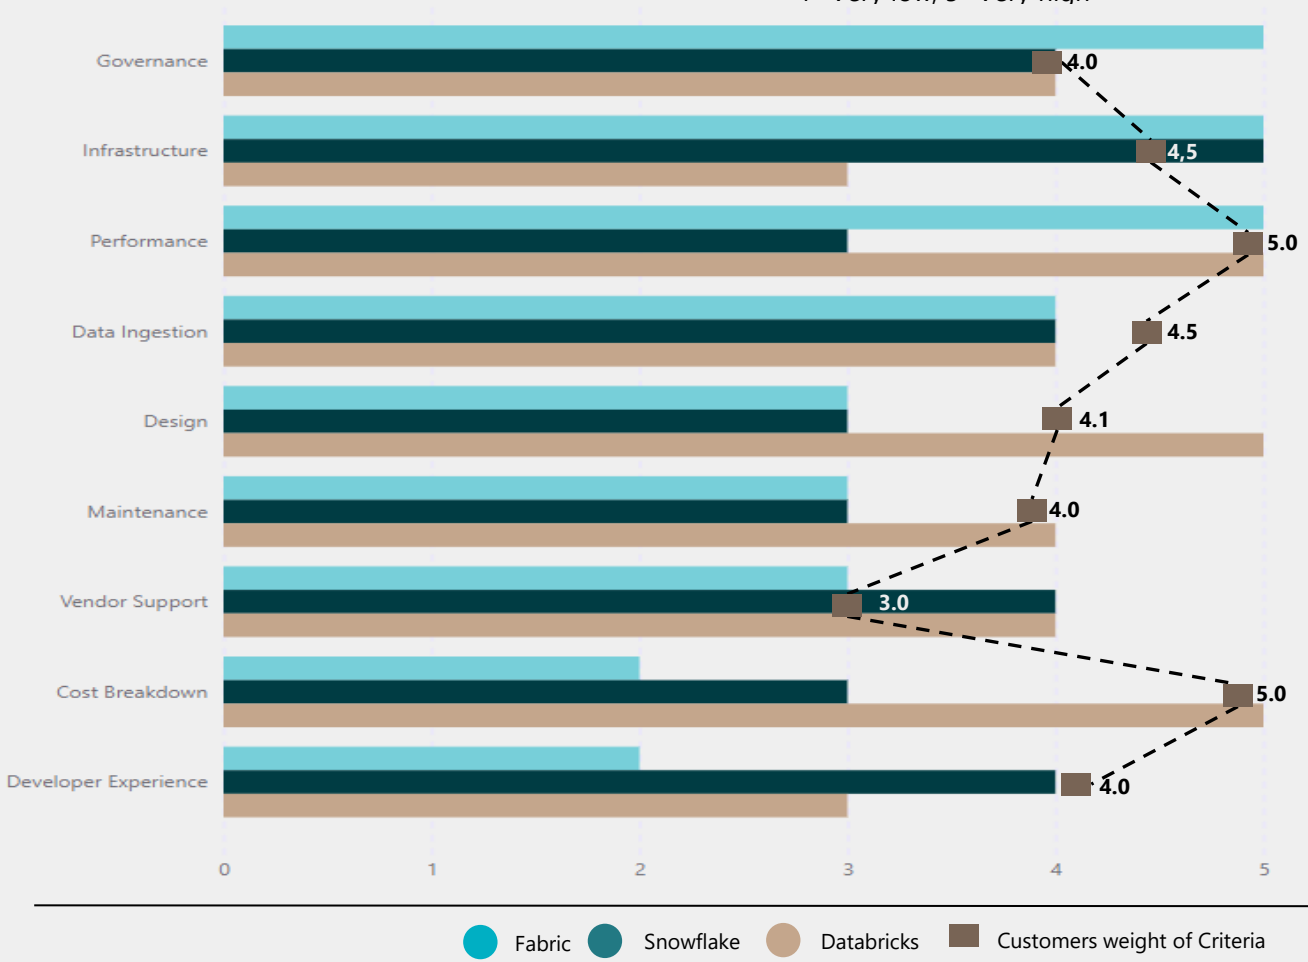

Recommended Platform

High-performance & Lowest cost
Databricks scores high on the most important customer criteria and is the recommended future data platform

Performance & Governance
Microsoft Fabric scored high on performance on the selected use case and scored highest on Governance

Usability & Vendor Support
Snowflake scores the highest while other platforms scores consistent with customer expectations

End-to-end
All platforms serves the end-to-end purpose required of the selected use cases

example

Cost Models

TCO	Higher complexity, improving with serverless	Lower complexity, but compute needs to be guardrailed	Depends on capacity sizing
Flexibility	Highly flexible (PAYG, reserved)	PAYG + capacity contracts	Fixed capacity model and reserve CU's
Transparency	Simple with serverless, complex otherwise	Mostly clear, some hidden costs	Very transparent
Monitoring	Requires active management	Strong tools, requires governance	Limited granularity

Developer Experience & Governance

			
Learning Curve	Moderate (SQL + Python/Spark knowledge beneficial)	Low (SQL-first, intuitive and easy to adopt)	Moderate (familiar tools, but evolving platform)
Workflow & Orchestration	Flexible, code-first workflows with strong CI/CD support & SDP with dependencies	Simple, well-documented and tasks & Dynamic Tables (simpler, less flexible)	Strong UI options, flexible design, less mature CI/CD support
Local Development	full support (Databricks connect etc)	Seamless with tools like VS Code and lightweight setup	Limited (most workloads require cloud execution – acknowledged)
Data Governance & Lineage	Strong lineage (table + column level) via Unity Catalog	Object-level lineage within platform (tables/views)	Requires Purview for full lineage and governance

Governance & Operational

Data Governance & Lineage	Strong lineage (table + column level) via Unity Catalog	Object-level lineage within platform (tables/views)	Requires Purview for full lineage and governance
Orchestration & Workflows	Advanced Jobs, SDP with dependencies and scheduling	Tasks & Dynamic Tables (simpler, less flexible)	Integrated pipelines (Data Factory + Airflow)
Infrastructure & Deployment	Full IaC support (Terraform), flexible cluster configuration	Managed service, strong Terraform support, minimal setup	Limited IaC (evolving), tightly integrated with Azure
Platform Control	Multi-cloud, high flexibility, requires active management	Multi-cloud, Less control, but simpler operations and governance	Microsoft-cloud, capacity-based model, centralized but less flexible

Agenda

- What is a modern data platform
- Platform architecture and overview
- Comparison of platforms
- AI-ready platforms and best of breed

AI-Ready Data Platform – Prerequisites

01

Own it

If nobody owns it, nobody fixes it. Clear ownership by domain, with a catalog that makes definitions and responsibility visible.

02

Measure it

If no one measures it, no one trusts it. Continuous quality rules – completeness, freshness, conformance, not a one-time cleanup.

03

Control it

Who can see what, and why? Access control down to column and row level, declared once, enforced everywhere.

04

Prove it

Where did this number come from? Automated lineage from source to report – auditable, explainable, regulator-ready.

AI Capabilities in Modern Platforms

Genie (NLQ on governed data)

Genie Code (agentic dev + workflows)

Agent Bricks (prebuilt AI agents)

Mosaic AI (Vector Search, Model Serving)

AI Gateway + MLflow (governance & monitoring)

Quick read:

Strongest for AI engineering & agent development (builder-focused).

Fabric Data Agents (NLQ across lakehouse/semantic models)

Fabric IQ (ontology + business context layer)

Operations Agent (real-time decisioning)

AI Functions (LLM in data pipelines)

Copilot across workloads.

Quick read:

Best for business-driven AI with semantic context + M365 integration.

Cortex Analyst (NL → SQL)

Cortex Agents (orchestration layer)

Cortex Search (RAG over data)

Cortex Code (AI dev agent – available in UI and CLI)

Snowflake Intelligence (end-user AI layer)

Quick read:

Most unified AI stack (SQL → agents → business users).

Why Platform Lock-In Is Risky

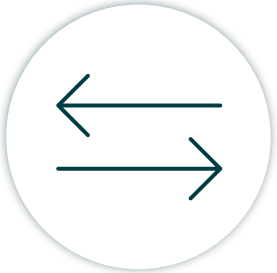

Platform changes

Technology evolves

Business needs change

Consider the Ideal Solutions

All-in one

A unified platform that integrates all the essential components of a modern data platform - from ingestion and transformation to orchestration - providing everything you need in a single streamlined solution.

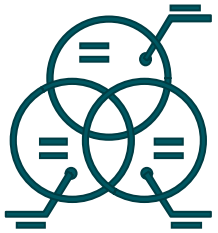

Best of breed

A tailored combination of tools designed to deliver an exceptional developer experience, leveraging the latest technologies and approaches to meet your specific needs.

Platform-Agnostic Data Platform Principles

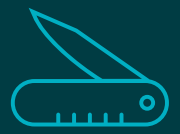

Use the right tool for the job, not just one platform

Pick tools based on needs (performance, cost, governance, dev experience)

Avoid vendor lock-in where it does not add value

Prefer tools that work well across platform

Let each tool do what it is best at

Best of breed: Mix & match

Data Catalog & Governance

OpenMetaData

Data Ingestion

dlt

Data Transformation

dbt

Data Orchestration

Dagster

Data Observability

Grafana

ARCHITECTURE

Two Platform Example

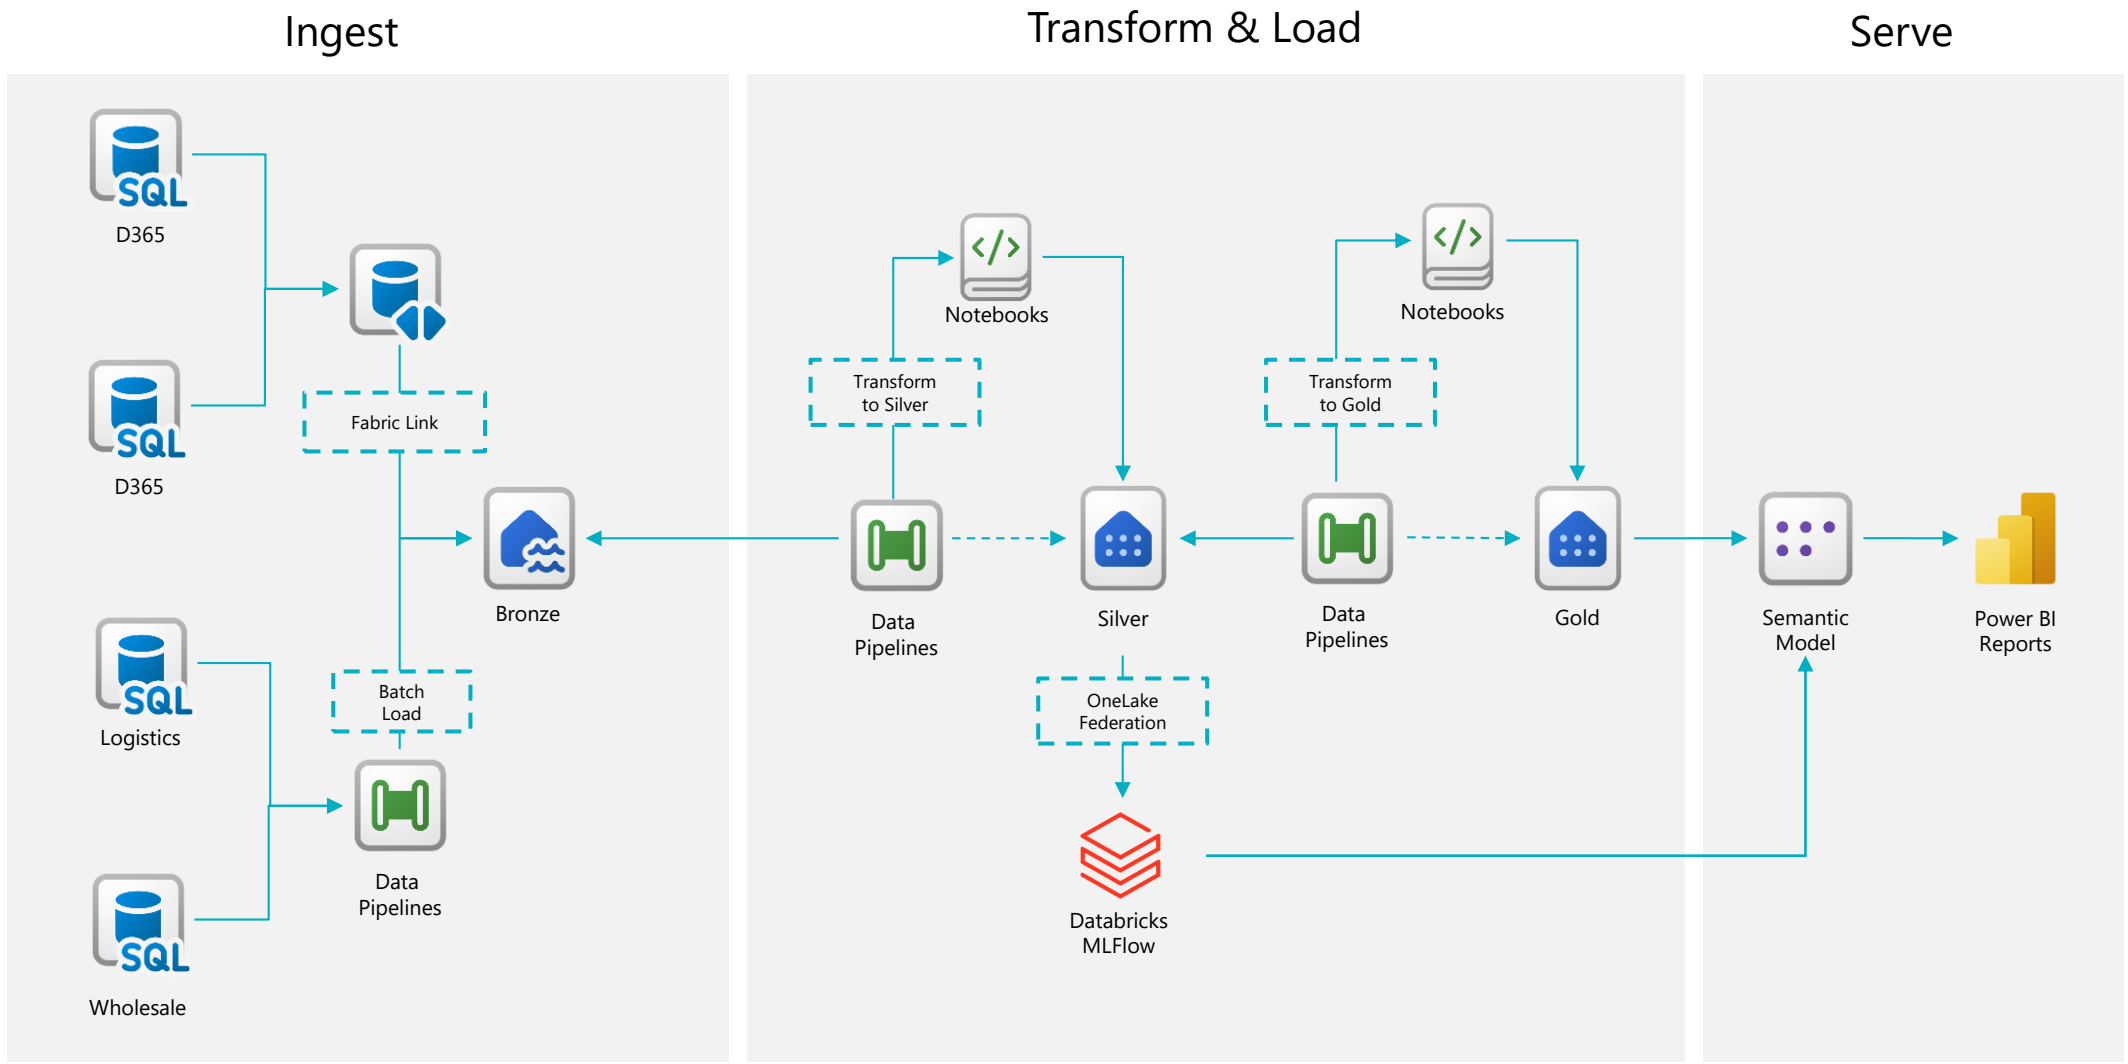

ARCHITECTURE

Best of Breed Example

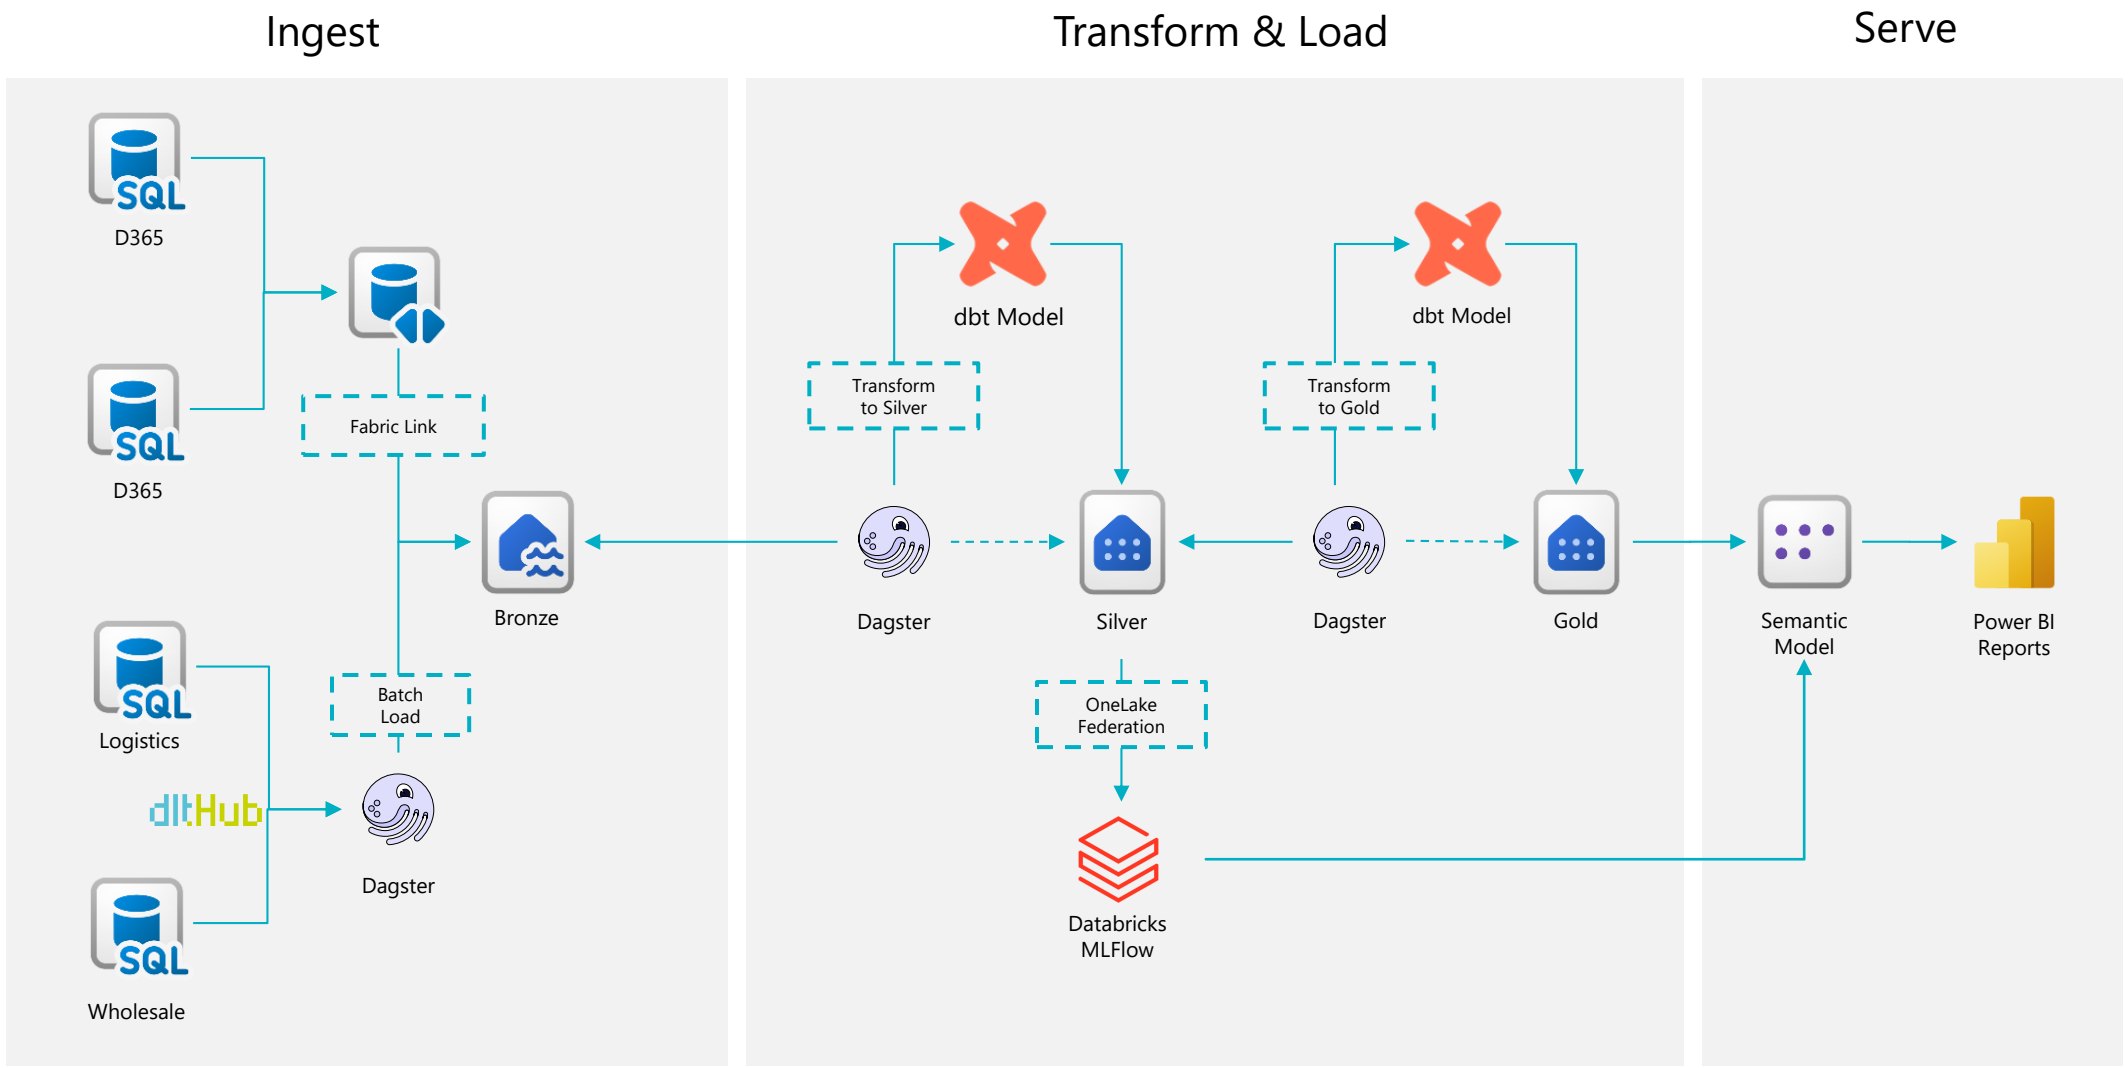

Key Takeaways

The landscape is complex

Platforms have different strengths

AI-ready foundation beats tools

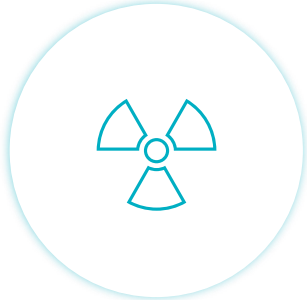

Platform-agnostic reduces risk

Time for Q&A

Daniel Størup Thomsen

dtho@inspari.dk

Mads Klavsen

mkla@inspari.dk

Get in touch!

In case you have questions, reflections, or seek sparring on something specific, please don't hesitate to reach out.