



2026 Trends in Analytic Architectures

Presented by: William McKnight

“#1 Global Thought Leader and Influencer in Cloud and in Data Center” Thinkers360

President, McKnight Consulting Group
3 x **Inc 5000**

 /in/wmcknight

www.mcknightcg.com
(214) 514-1444



McKnight Consulting Group Partial Technology Implementation Expertise

Big/Analytic/Vector/Mixed Data Management



Data Movement and APIs

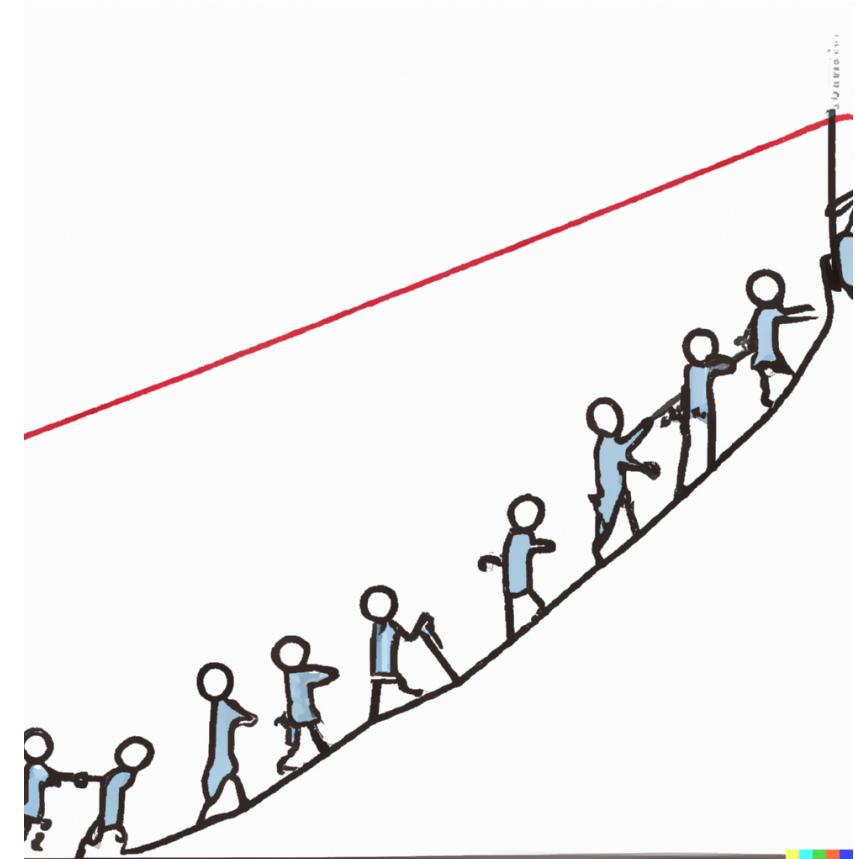


Operational/Transactional Data Management



Why Are Trends Important?

- It is imperative to see trends that affect your business to know how to respond
- Plan for and deal with change
- Better to be at the beginning of the trend rather than the end
- Wants, needs, and tastes of your customer changes
- Make you a leader, not a follower
- Grow your business ideas
- Give you ideas what to improve in your business



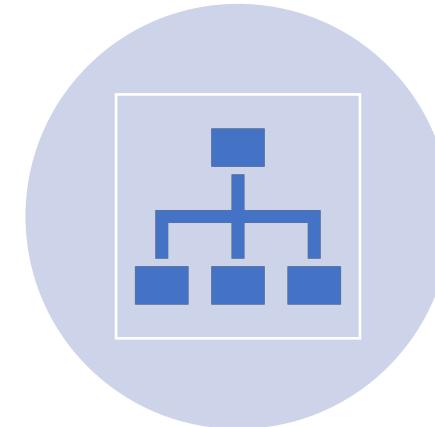
Information Management Leaders



INFORMATION MANAGEMENT
LEADERS OF TOMORROW CAN
ADVANCE MATURITY WHILE
ALSO SOLVING BUSINESS
ISSUES



THERE'S NO BUDGET FOR
“STAYING ON TRENDS”



INFORMATION MANAGEMENT
LEADERS MUST PICK THEIR
WINNING (I.E., MULTI-YEAR
SUSTAINABLE) APPROACHES
AND GET ON BOARD

2025

- **Agentic AI hit reality:** Widespread experimentation exposed limits in reasoning, reliability, and memory, slowing large-scale production impact.
- **Standards emerged fast:** MCP quickly became the de facto way models connect to data, APIs, and agent workflows.
- **Data foundations consolidated:** PostgreSQL, Zero ETL, and data sharing became core patterns for unifying transactional and analytical systems.
- **Open architectures won:** Apache Iceberg solidified its role as the standard open table format across the ecosystem.
- **AI drove real economic impact:** Productivity gains from AI accounted for the majority of U.S. GDP growth, signaling a shift from pilots to material value creation.
- Renewed interest in data
- Insane Capital Raises: Examples include OpenAI raising a staggering \$40 billion, reflecting an exponential growth trajectory.

Last Year's Trends that Hit

Information Management Trends

- Simplified data architecture
- Consolidation, rationalization & backend optimization
- Optimization by reducing redundant/unused data
- Open table formats (Iceberg, Delta, Hudi) widely adopted
- Data privacy impacting business operations
- Regulatory compliance costs rising
- Data lakehouses becoming the de facto standard
- Unstructured data at parity with structured data
- Data mesh simplifying governance/security
- Data products as standard exchange mechanism
- EU AI Act impact

Artificial Intelligence Trends

- MLOps adoption
- Executives expecting natural language with data
- AI challenges (explainability, bias, hallucinations, governance)
- 2025 as the Year of AI Agents
- AI at edge
- AI companions/replicas

Broader Environment Trends

- Cloud-native principles (DevOps, containers, microservices)
- Data observability gaining ground
- Education & healthcare adoption of AI rising
- Entry-level jobs becoming harder to find due to AI
- Hybrid quantum computing investment growing
- Data engineering as highest-valued profession
- 2025 as the year of data observability
- Entry-level jobs universally harder to find

Last Year's Trends that Missed

Information Management Trends

- Unstructured data equal to structured in tooling & parity

Artificial Intelligence Trends

- AI agents pervasive for complex, specific tasks
- GenAI & LLM success widely evident
- Pervasive adoption of AI agents for complex tasks
- Neuromorphic/optical computing advancements
- Executives fully trusting autonomous AI decisions

Broader Environment Trends

- Hybrid quantum computing delivering exponential big-data query performance

The Goal of Modern Enterprises

- **Automating complex tasks:**
enabling unprecedented levels of efficiency and productivity.
- **Providing smart data environments:**
offering personalized assistance for data exploration and insights.
- **Driving proactive decision-making:**
unlocking new levels of data insights and automation.
- **Simplifying business processes:**
streamlining operations and improving ROI.
- **Facilitating autonomous decision-making:**
enabling hyper-personalization and data-driven decision-making.

The World is Changing

● Old Way:

Ramp up: Spend months getting familiar with the company

Data discovery: Take hours to find and understand relevant data

Waiting on data engineers: Spend hours to days waiting for data engineers to ingest data

Data optimization: Wait days to weeks for data engineers to optimize data

Dashboard creation: Spend days to weeks building business dashboards and spaces

● New Way:

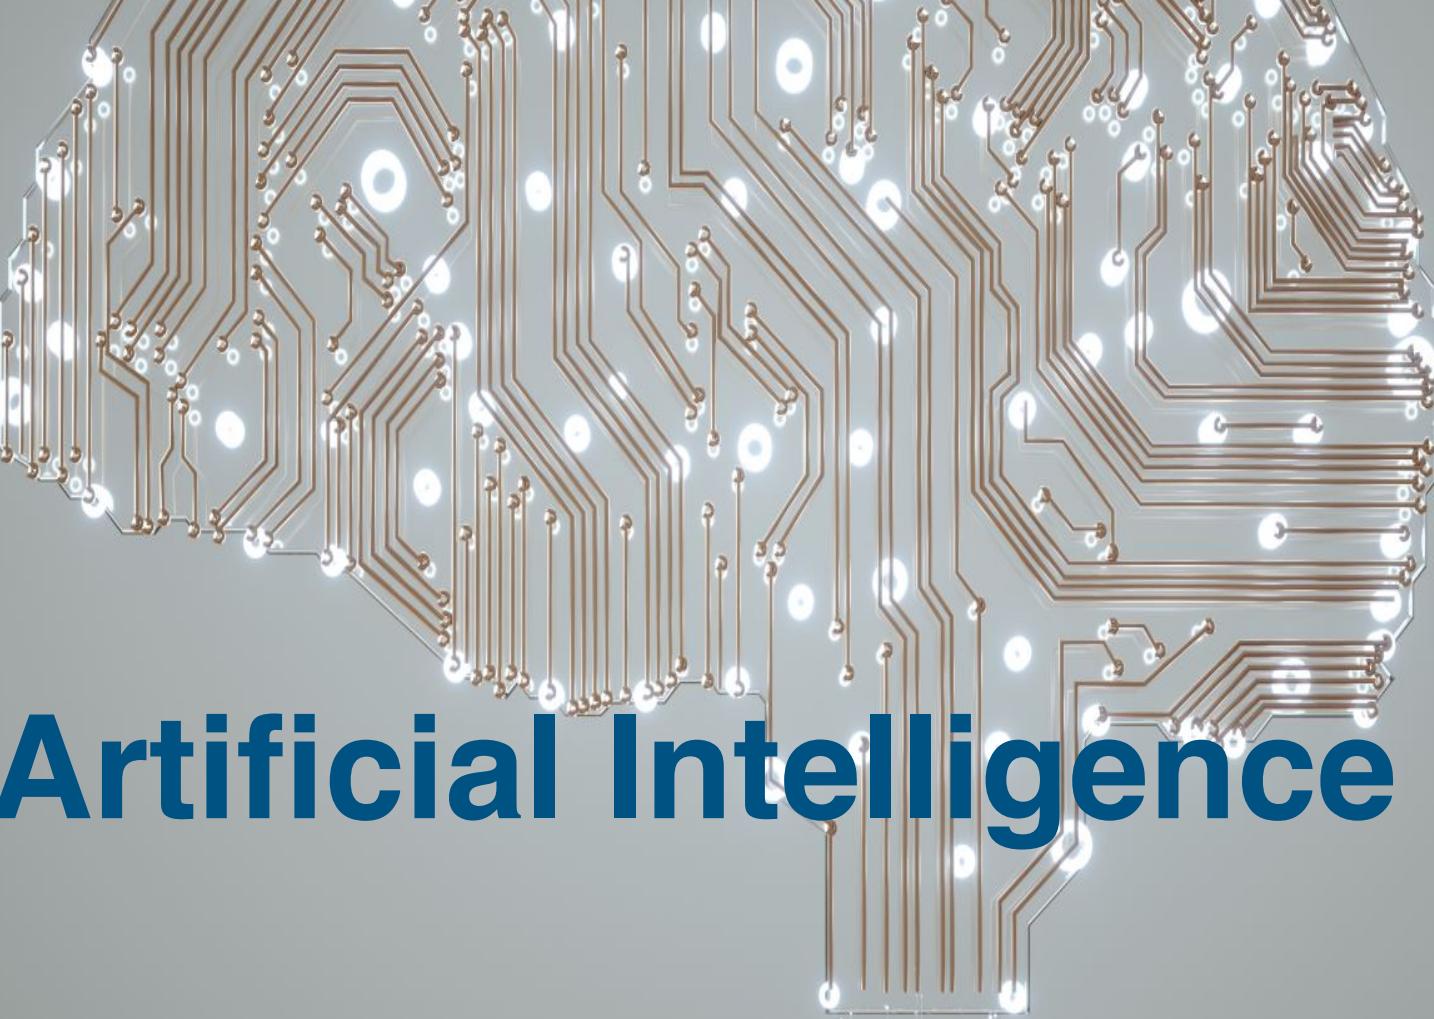
Instant insights: AI agents interface with structured and unstructured data

Smart navigation: Semantic layer guides agents to the data instantly, providing business context

Unified access: Zero-ETL federation gives unified access to data instantly

Fast answers: Agents deliver answers instantly

Optimized queries: Agents optimize queries for everyone instantly



Artificial Intelligence

The AI Adoption Curve: 2025–2026

2025 focused on pilots and partial automation

Adoption still on flat hockey-stick blade

2026 marks rapid acceleration

Autonomous execution scales

Investment velocity increases



Best of Breed AI Architectures



Fungible AI Stacks



Smaller, More Efficient LLMs



Rapid model obsolescence as a strategic risk



Economics over raw capability

AI Agents vs Generative AI

GenAI focused on content generation

AI agents execute multi-step tasks

Agents combine models, data, and software

Workflow orchestration across systems

Agents deliver higher enterprise ROI

From Generative AI to Agentic AI

Shift from prompt-based GenAI to autonomous AI agents

Agents reason over context and enterprise data independently

Metadata-aware knowledge layers critical for trust and accuracy

Agents surface insights without human prompting

Used for contracts, fraud detection, and hyper-personalization

On-Prem AI Infrastructure

Rising demand for data sovereignty and compliance

On-prem RAG platforms reduce development complexity

Lower complexity vs DIY cloud stacks

Lower total cost of ownership

Mission-critical AI repatriated by 2026

Sovereign AI

Integrated RAG platforms outperform DIY stacks

Lower development complexity

Lower ongoing maintenance effort

Faster time to production

Critical for agentic AI pipelines

Quantified ROI

Automated customer service ~400% ROI

High autonomous resolution rates

Fraud detection preventing major losses

Fraud ROI exceeding 400%

MDM consolidation delivering recurring savings

Natural Language as the Primary Interface



Natural language as the default control layer



Lowering the skill barrier



Compression of complex workflow



Context-aware interaction



Abstraction of underlying complexity



Speed of iteration and experimentation



Strategic control point

Semantic Layers as Foundation

Semantic layers
as the connective
tissue

Decoupling logic
from
infrastructure

Foundation for
interoperability

Consistency and
governance at
scale

Acceleration of AI
deployment

Improved
reasoning and
reliability

Long-term
strategic asset

2026 Outlook: The Digital Researcher

- GenAI as digital librarian phase ending
- Agentic AI acts as digital researcher
- Reads, reasons, and connects enterprise data
- Produces insights and decisions autonomously
- Marks inflection point in AI value creation



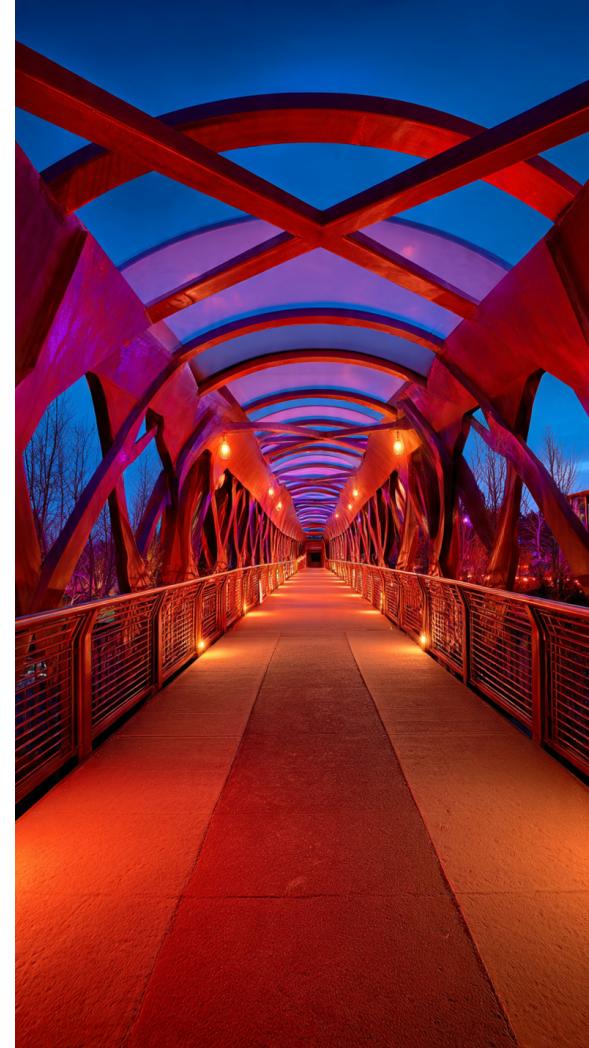
The Technical Divide

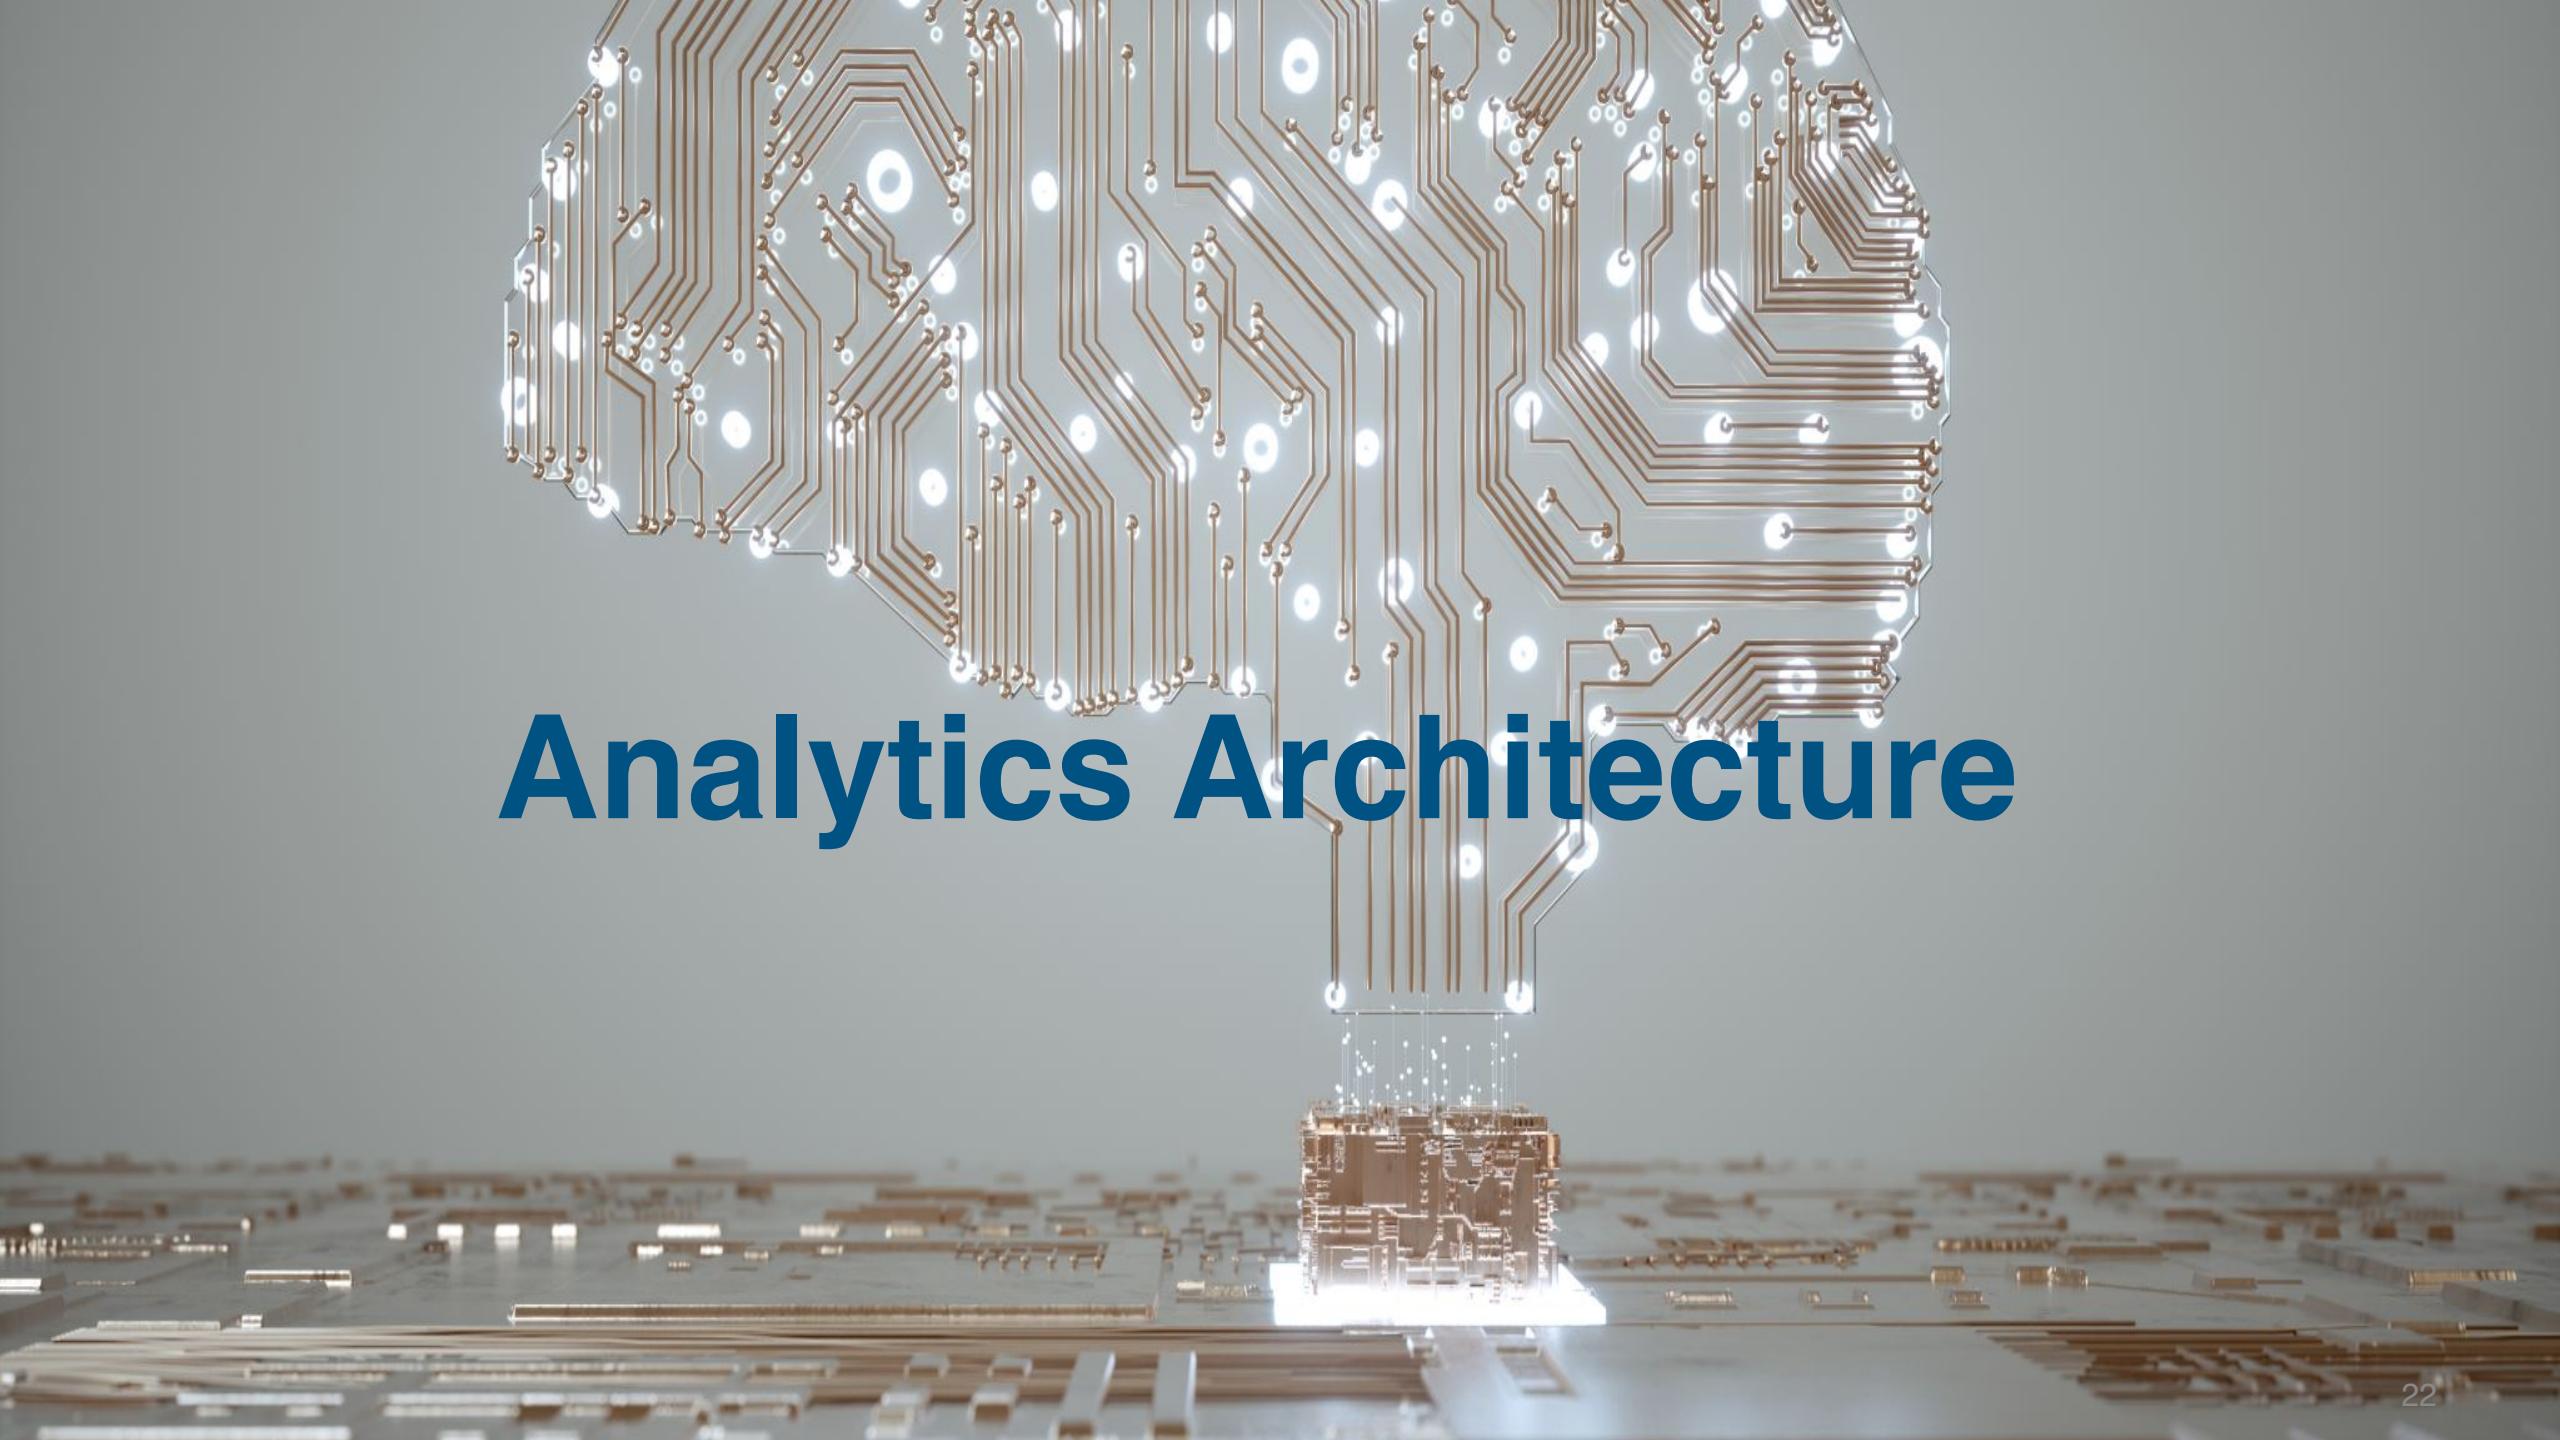
AI is leveling the field, enabling non-technical users to perform tasks once reserved for specialists; Enterprise users use AI for work they previously could not do

Advantage is shifting from pure technical skill to domain expertise paired with AI execution

Value is moving from prompt wording to providing rich business context, since models lack awareness of internal goals, data, and rules

Platforms are embedding AI to capture this context, creating lock-in as AI improves with deeper data access





Analytics Architecture

Modern Data Architectures



Icehouses: ~80% savings on incremental loads

Columnar formats accelerate analytics

Faster joins and scans

Open architectures reduce lock-in

Automated Master Data Management

LLMs generate and enrich master data

AI-driven entity resolution

NER and NLU mapping records automatically

Systems learn from stewardship decisions

Humans focus on exceptions

Data Fabric & Semantic Layer ROI

Business data fabric reduces cognitive load

63% lower initial build cost vs DIY

72.4% reduction in data modeling effort

2.4x more cost-effective over three years

Semantic layers accelerate decision-making

Specialized Hardware & Edge AI



Real-time inference moving to edge

Specialized low-power AI chips

Embedded zero-DBA databases

Zero-ETL cloud synchronization

Autonomous factories emerging

Unified & Multi-Model Platforms

Move away from stitched-together point solutions

Consolidation reduces data debt and operational complexity

Single engines handling transactional, analytical, and NoSQL workloads

Predictable cost and simpler administration

Multi-model databases favored by 2026

Open Data Lake Format Dominance

Open lake formats becoming AI and analytics foundation

Iceberg and Delta Lake enable compute–storage separation

77–95% lower compute and ingestion cost vs warehouses

Interoperability reduces vendor lock-in

Iceberg emerging as long-term storage standard



Python-Centric Real-Time Ecosystems

Developer preference shifting toward Python-first frameworks

Lower memory usage than Java-based alternatives

Up to 4.6x more cost-effective than Flink

Faster iteration and development cycles

AI apps reach production 3–9x faster

Distributed SQL & Transactional Performance

- Cloud-managed distributed SQL validated at scale
- Unified platforms beat by 14x price-performance
- Lower 3-year TCO for transactional workloads



Digital Analytics & Behavioral Data

Shift from vendor UI to warehouse-first analytics

Using an event data platform with transformation-based accelerators results in queries that are significantly simpler and dramatically faster

42% less complex queries

6.9x faster analytics performance

Over 99% reduction in data scanned



Foundation Concepts

Lakehouse = open tables + compute separation

Microservices are independently deployable

APIs, containers, orchestration, observability

RAG grounds LLMs in specific knowledge

Maturing Data Architectures Misc.

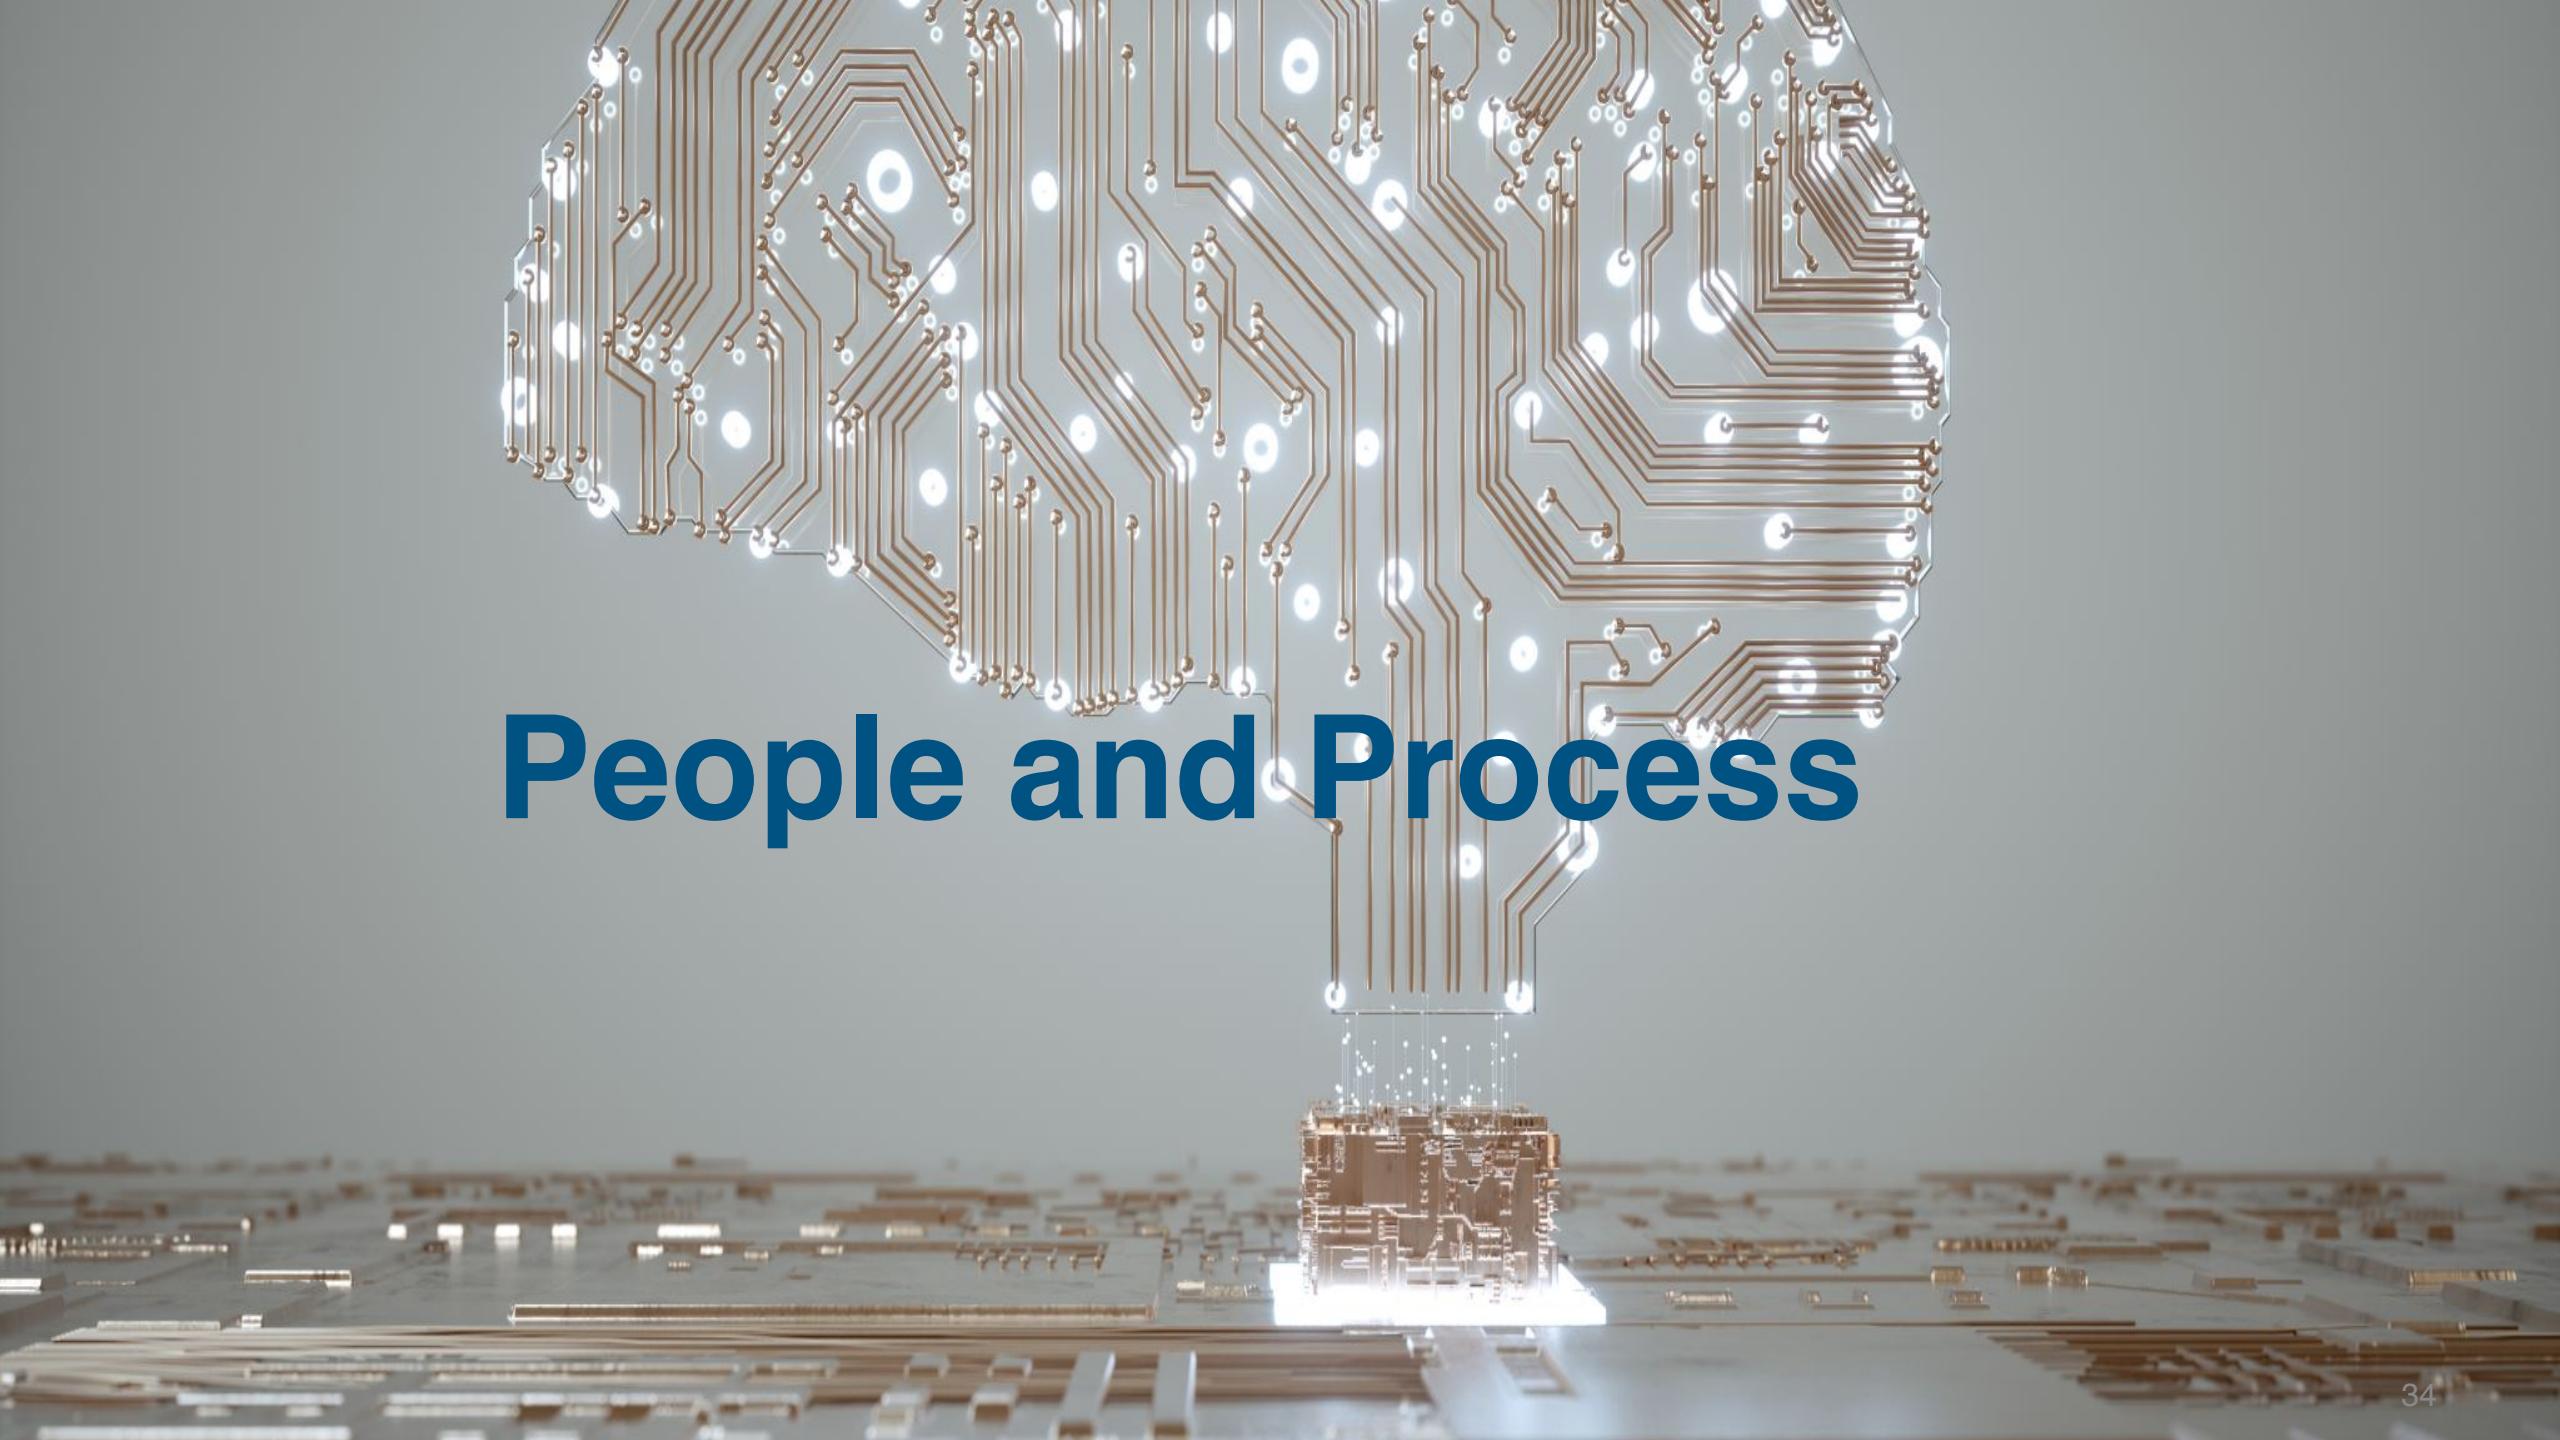
Unstructured data nearing parity with structured data

Text, audio, and video driving competitive advantage

Data mesh moving from theory to execution

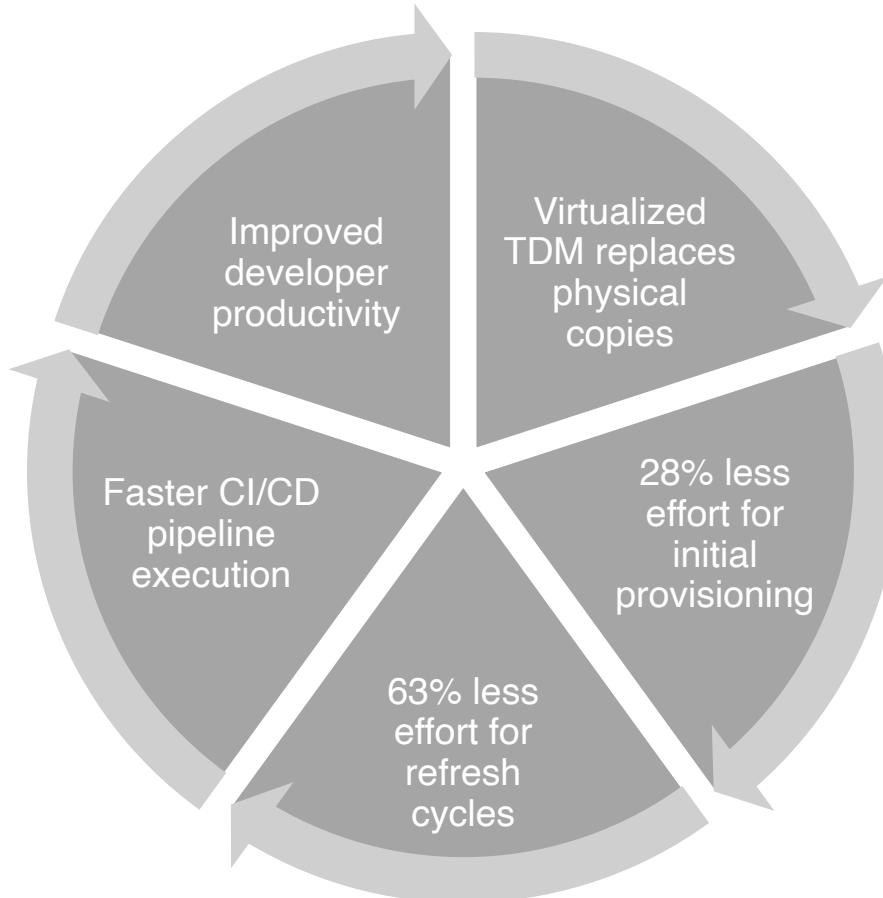
Decentralized ownership with shared standards

Simplified governance and security models



People and Process

Test Data Management & DevOps Velocity



The Rise of Autonomous Digital Operations

Accelerating AIOps spend across startups and mid-market

Shift from AI-assisted monitoring to autonomous remediation

Self-healing platforms resolving incidents automatically

Agentic AI orchestrating end-to-end workflows

Autonomous execution in use by 2026

Advanced Computing & Financial Shifts



Early hybrid quantum traction in 2025

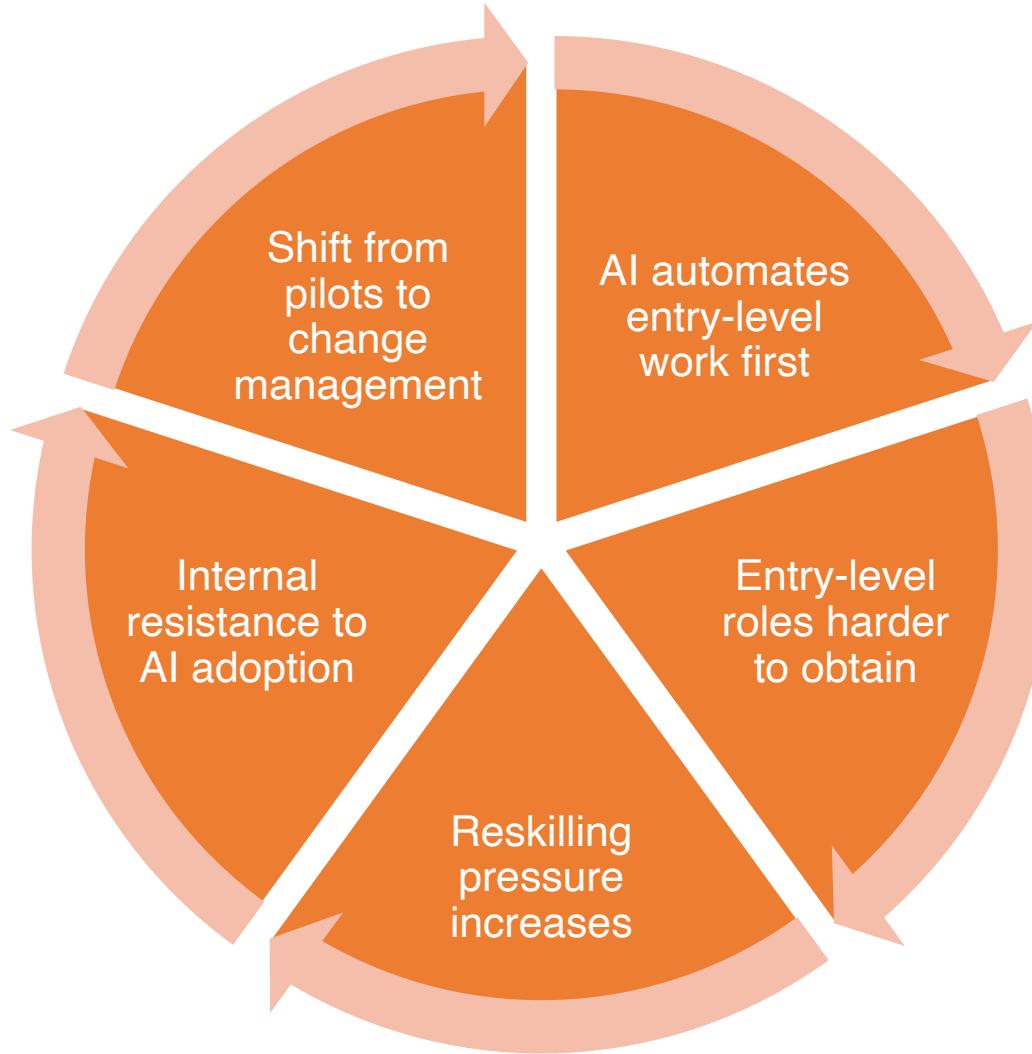
2026 as quantum investment inflection point

Data FinOps becomes standard discipline

Backend rationalization accelerates

Cost-tiered storage and compute strategies

Social & Economic Impact



Emerging Enterprise Risks



- AI-generated data exceeds half of new content
 - Risk of model collapse
 - Data poisoning attacks
 - Need for lineage and validation
 - Expanded AI governance requirements

Bring Your Own Cloud (BYOC) Models

Software deployed inside customer cloud perimeter

Up to 60% cheaper than SaaS streaming platforms

Improved data sovereignty and security

Leverages reserved instances and savings plans

Managed experience without SaaS lock-in



Cloud Marketplace Procurement Agility

Marketplace-based procurement accelerates deployment

Licensing reduced from weeks to minutes

Improved time-to-value

Standardized purchasing and billing

Greater flexibility in enterprise buying

Evolution of Professional Roles

- AI architects essential for enterprise AI maturity
- Data engineering remains highest-value profession
- Natural-language access to data becomes table stakes
- Rise of AI ethicists, auditors, and trainers
- Trust, governance, and domain expertise prioritized



Looking Ahead to 2026

Autonomy becomes competitive differentiator

Agentic AI + open data converge

Governance embedded by design

Leaders separate from laggards

Execution replaces experimentation

Summary

Agentic AI

- AI moves from standalone generative models to **agentic systems** that execute tasks, orchestrate workflows, and drive measurable business outcomes, with cost, latency, and scale now outweighing raw model power.

Semantic Layers

- **Semantic layers and governed data foundations** become essential, enabling interoperability, consistent meaning, faster AI deployment, and long-term value beyond individual models.

Data Architecture

- Data architectures will converge on **open formats, distributed SQL, real-time pipelines, and automated master data management** to support scalable AI decision-making.

Sovereign AI

- Infrastructure will split between **sovereign/on-prem** and **edge** for control and low latency, and cloud/**BYOC**/**marketplace** for flexibility and speed.

Organizational

- Organizations will adapt roles, operating models, **DevOps**, and governance as **autonomous** systems increase productivity, **ROI**, and shift human focus to **oversight, exceptions, and strategy**.



2026 Trends in Analytic Architectures

Presented by: William McKnight

“#1 Global Thought Leader and Influencer in Cloud and in Data Center” Thinkers360

President, McKnight Consulting Group
3 x **Inc 5000**

 /in/wmcknight

www.mcknightcg.com
(214) 514-1444

