AI-ready data architecture

A modern, governed and scalable data foundation that provides trusted, high-quality and context-rich data to power robust, enterprise-grade AI workloads

EY
Shape the future with confidence

The better the question. The better the answer. The better the world works.

Contents

1. Executive summary

Organizations across industries are accelerating the adoption of artificial intelligence (AI) to drive operational efficiency, unlock new revenue opportunities and deliver increasingly personalized customer experiences. For senior data and technology leaders operating in complex and regulated environments (CDOs, CTOs, IT leaders, and enterprise or data architects), the priority has shifted from experimentation to the responsible scaling of AI. However, the ability to realize value at scale is fundamentally dependent on the availability of high-quality, well-governed and readily accessible data.

Traditional data architectures, characterized by siloed systems, batch-oriented processing and limited scalability, are not designed to meet the requirements of modern AI and generative AI (GenAI) workloads. These workloads demand real-time access to data, consistent business semantics and the ability to manage high volumes of structured and unstructured data with embedded governance and control. In the absence of these capabilities, AI initiatives often remain confined to pilots or fail to deliver reliable, production-grade outcomes.

An AI-ready data architecture provides a governed, modern and scalable foundation designed specifically to support AI at enterprise scale. It enables continuous access to trusted, interoperable and context-rich data, while embedding governance, security and data quality controls across the entire data lifecycle. By integrating cloud-native services, real-time data pipelines, unified governance frameworks, and semantic and vector data capabilities, it supports advanced analytics, machine learning (ML), large language models (LLMs) and intelligent automation.

Importantly, AI-ready data architecture redefines governance as an integral, embedded capability rather than a reactive control function. It ensures that data remains trustworthy, compliant and fit for purpose across both human and machine-driven consumption, thereby reducing operational risk and enabling consistent, scalable AI adoption.

This whitepaper outlines the AI-ready data conceptual architecture pattern, key principles and modern AI data supply chain capabilities required to establish an AI-ready data architecture. It is intended to support organizations in defining modernization priorities, establishing enterprise standards and building a resilient data foundation that enables sustainable, AI-driven transformation.

2. Why AI-ready data architecture matters

As organizations accelerate their adoption of artificial intelligence (AI), the focus is rapidly shifting from experimentation to enterprise-scale deployment. Yet most traditional data architectures, built for reporting and transactional workloads, are fundamentally misaligned with the demands of modern AI and GenAI. This misalignment creates a critical gap between ambition and execution, limiting the ability to operationalize AI at scale.

AI systems are only as effective as the data that powers them. They require continuous access to high-quality, context-rich and governed data. In the absence of this foundation, initiatives stall — driven by inconsistent definitions, limited real-time data access and fragmented governance. As a result, organizations remain trapped in pilot cycles, unable to deliver trusted, production-grade AI outcomes.

An AI-ready data architecture directly addresses these challenges. It provides a governed, modern and scalable foundation purpose-built for AI workloads, enabling real-time data access, enforcing unified governance and ensuring that data remains secure, compliant and trustworthy for both human and machine consumption. This is not an incremental improvement; it is a necessary shift to support AI-driven operations.

This shift builds on the evolution of data platforms, from warehouses to lakes and lakehouse architectures, but goes further to meet AI-specific demands. While earlier architectures improved analytics, they fall short in areas critical to AI, including semantic consistency, real-time processing and governance of unstructured and multimodal data. AI-ready architecture closes this gap by integrating these capabilities with centralized policy orchestration — often through a control tower model — to standardize governance, automate controls and enable secure, scalable data access across the enterprise.

Ultimately, the move to an AI-ready data architecture is not a technology decision; it is a strategic one. Organizations that invest now can accelerate time-to-value, strengthen trust in AI outcomes and reduce operational and regulatory risk. Those that do not risk falling behind. AI-ready data architecture is the foundation for scaling AI with confidence, delivering faster insights, better decisions and sustained competitive advantage in an AI-driven economy.

3. Challenges in scaling AI with current data foundations

Despite rapid advancements in cloud platforms and AI tooling, organizations still face significant architectural and operational gaps when attempting to operationalize AI at scale. The core issue is not the lack of AI tools, but data foundations designed for reporting and transactions rather than continuous AI consumption. Traditional data architectures, built primarily for reporting and transactional systems, struggle to support the velocity, variety and governance demands of modern AI systems.

Siloed and fragmented data ecosystems

Enterprise data is often distributed across multiple warehouses, data lakes, legacy on-premises platforms, software as a service (SaaS) applications, and streaming systems. This fragmentation results in:

- Duplication and inconsistent data definitions
- Delayed data availability for AI model training
- Complex, brittle integration pipelines

Impact: **AI teams make significant effort reconciling data instead of building and scaling models, slowing delivery and limiting reuse.**

Lack of real-time data availability

Most legacy architectures rely heavily on batch processing and manual data pipelines.
As a result:

- AI models operate on stale or delayed data sets.
- Real-time use cases, such as personalization, fraud detection, or dynamic optimization, become impractical.

Impact: **AI remains largely offline or experimental, unable to support real-time decision-making in production environments.**

Insufficient high-quality data and metadata

AI outcomes are only as reliable as the data sets that power them. Many organizations struggle with:

- Missing, inconsistent or outdated records
- Limited automation for data quality validation
- Incomplete lineage and traceability
- Manual, error-prone transformations

Impact: **Model development slows, trust in predictions declines, and auditability becomes harder to sustain.**

Absence of shared ontology and semantic alignment

Without a unified semantic layer or enterprise ontology:

- Business concepts are defined inconsistently across data systems.
- Interoperability between data sets is reduced.
- AI models receive incomplete or conflicting context.

Impact: **Teams duplicate effort, and scaling AI across domains becomes challenging.**

Lack of unified governance, security and privacy controls

AI workloads frequently require sensitive and regulated data, yet many environments lack:

- Fine-grained access control
- Robust masking and tokenization
- Automated enforcement of compliance requirements (e.g., General Data Protection Regulation (GDPR), Health Insurance Portability and Accountability Act (HIPAA), data residency)
- Monitoring to detect unauthorized or inappropriate data usage

Impact: Regulatory, ethical and operational risks increase, constraining broader AI adoption.

Inadequate support for unstructured data

Modern AI systems demand access to text, documents, images, logs, audio and sensor data. Traditional business intelligence (BI)-focused architectures typically lack:

- Efficient storage for large unstructured data sets
- Scalable indexing and search capabilities
- Tools to process and analyze unstructured data at scale

Impact: Enterprises struggle to enable LLMs and multimodal AI solutions.

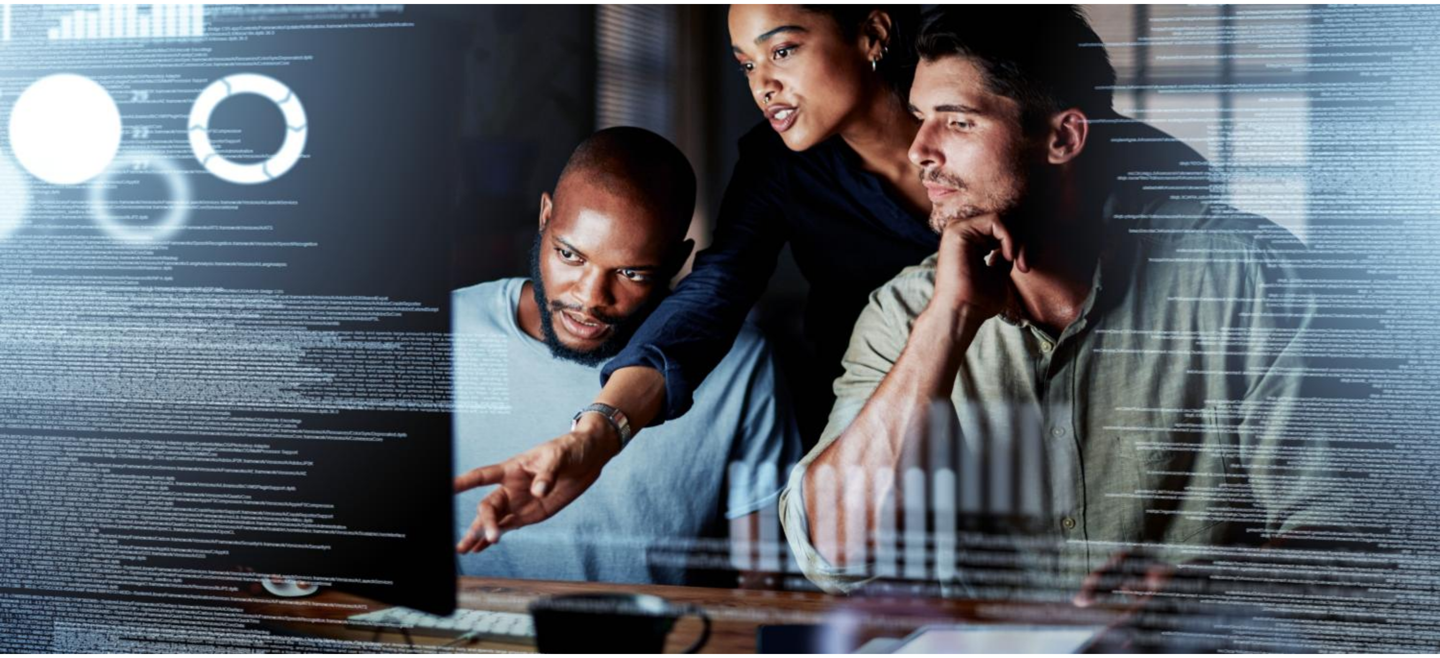

These challenges persist largely because organizations attempt incremental modernization of architectures that were never designed for AI. Overcoming them requires a new, AI-ready architectural foundation.

4. Evolution to AI-ready data architecture

Enterprise data architectures have evolved significantly over time, reflecting the growing complexity of data and the increasing demands of business and technology. This progression, from traditional data warehouses to data lakes, lakehouse architectures and ultimately AI-ready data architectures, highlights a shift from structured reporting toward enabling intelligent, real-time and scalable AI systems.

Traditional data warehouses were designed primarily for structured reporting and BI. They relied on a schema-on-write approach, enforcing strict data models and governance controls. While this provided strong data quality and consistency, it introduced rigidity, limited scalability and minimal support for advanced analytics or AI workloads. As a result, their suitability for modern AI use cases remains low.

The emergence of data lakes addressed some of these limitations by enabling storage of structured, semi-structured and unstructured data at scale using schema-on-read. This improved flexibility and supported experimentation with machine learning. However, weak governance, lack of semantic consistency and limited data quality controls reduced trust and constrained enterprise-wide adoption for critical AI use cases.

Lakehouse architectures evolved to bridge this gap by combining the reliability of data warehouses with the flexibility of data lakes.

They introduced hybrid schema approaches, improved governance and stronger support for analytics and machine learning workloads. While lakehouses significantly enhanced interoperability and operational balance, they still fall short of fully addressing the demands of real-time, large-scale and production-grade AI systems.

An AI-ready data architecture represents the next stage in this evolution. It is purpose-built for end-to-end AI, ML and GenAI enablement, supporting multimodal data, vector representations and advanced processing at scale. Unlike earlier architectures, it embeds governance, lineage and trust directly into the data life cycle, adopts semantic and ontology-driven models for consistent interpretation, and enables native real-time data pipelines. Interoperability is maximized through open standards, while operational efficiency is driven through automation and intelligent data management. As a result, AI-ready architectures provide very high suitability for enterprise-scale AI.

The shift toward AI-ready architecture is defined less by storage innovations and more by advancements in governance, standardized semantics and real-time data capabilities. These elements are critical for enabling trusted, scalable and explainable AI, particularly in environments where autonomous and agentic systems increasingly depend on consistent, high-quality data to operate effectively.

From warehouse to AI-ready: understanding architectural advancements

Dimension	Traditional data warehouse	Data lake	Lakehouse architecture	AI-ready data architecture
Primary purpose	Structured reporting, BI	Store raw data	Unified analytics + ML	End-to-end AI or ML or GenAI enablement
Data types supported	Structured	Structured + semi or unstructured	Structured + semi or unstructured	Structured, semi or unstructured, multimodal + vector
Schema approach	Schema-on-write	Schema-on-read	Hybrid	Semantic + ontology driven
Governance and quality	Strong but rigid	Weak	Improved	Embedded governance, lineage, trust
Real time and streaming	Limited	Supported	Supported	Native real-time pipelines
ML or AI support	Minimal	Good for experimentation	Strong	Production-scale ML or GenAI
Interoperability	Low	Medium	High	Very high (open standards)
Business meaning or semantics	Rigid	None	Limited	Rich semantic layer
Operational efficiency	Slow	High scale, low governance	Balanced	Max efficiency via automation
AI suitability	Low	Medium	High	Very high

5. AI-ready data conceptual architecture and data pattern

The AI-ready data conceptual architecture establishes a governed, end-to-end framework for creating, managing and consuming enterprise data products in a highly regulated and federated environment.

At its core is the control tower, which centralizes governance through policies, curated controls, and approval workflows to ensure compliance across the entire data lifecycle. Data is generated within domain-aligned data supply chains, where source data is transformed into reusable, policy-compliant data products. These products are then registered in the enterprise data catalog, the authoritative repository capturing metadata, lineage, and control attributes.

The semantic layer enhances data with shared business context, enabling consistent definitions, relationships across domains, and AI-ready understanding. Access is tightly governed through two layers: the metadata access layer, which controls what users can discover based on policies and context, and the data access layer, which enforces how authorized users interact with underlying data assets. All capabilities are exposed through a unified data marketplace, providing a policy-aware interface for discovery, access requests, and controlled data consumption at scale.

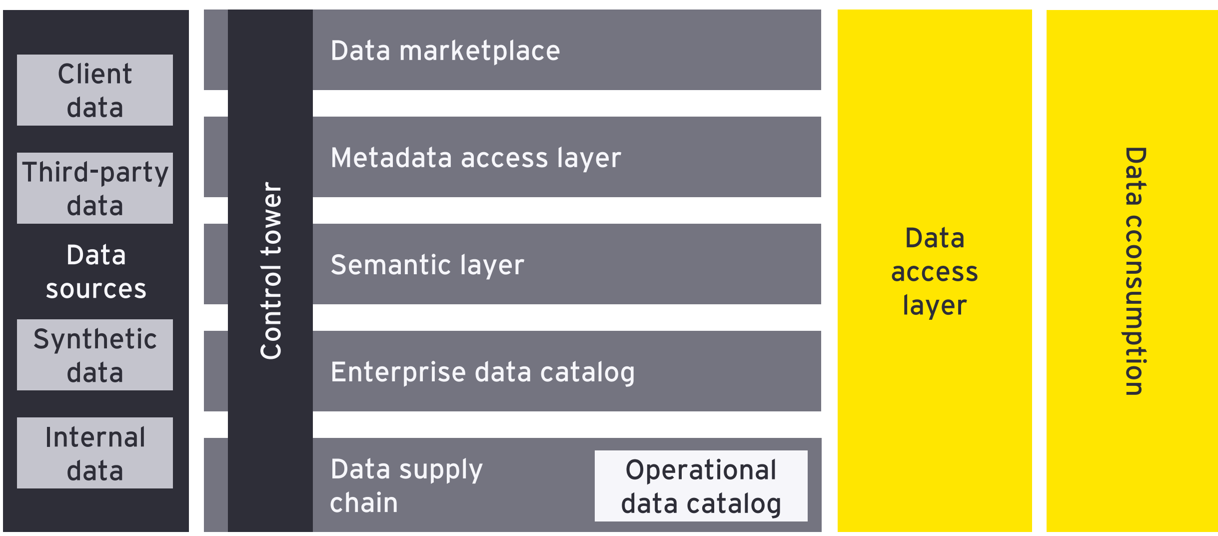

Figure 1 — AI-ready data conceptual architecture

This federated architecture enables secure data sharing without moving or duplicating data. Instead, governance is enforced centrally via metadata, policies, and access controls. Any data product available in the marketplace has undergone rigorous validation and remains continuously governed. Policy updates in the control tower propagate instantly across the ecosystem, ensuring enterprise-wide consistency without operational friction.

The platform promotes reuse by allowing business units to leverage certified data products without repeating governance efforts while still maintaining domain autonomy. Providers can use domain-specific tools as long as they integrate with platform standards. End-to-end traceability, from raw data to consumption, ensures transparency, with every data product carrying provenance details and all access being logged. This builds trust with stakeholders and regulators while enabling scalable, compliant data-driven innovation across the organization.

Component	Description
Control tower	The control tower provides centralized governance over data products by defining and enforcing enterprise-wide policies. It evaluates and enforces curated controls ("green lanes") throughout the data product lifecycle, from creation and publication to access, usage, and retirement. It automates checks (e.g., regulatory compliance, contractual obligations, privacy rules, data risk assessments) and coordinates human approvals when required. It provides a single point for policy definition and ensures every shared data product consistently adheres to corporate, legal, and client requirements.
Data supply chains	Data supply chains provide distributed pipelines for producing and managing data products. These include domain-specific platforms that create and host data products, each optimized for specific data types or business needs. These supply chains remain federated, allowing data to reside in source or appropriate platforms, but integrate with the control tower and catalog so that every data product – regardless of origin – is registered with required metadata and automatically governed by relevant controls and policies before being exposed in the marketplace.
Enterprise data catalog (catalog of catalogs)	The enterprise data catalog provides a unified view of metadata across data products. It aggregates metadata from distributed operational data catalogs and source systems into a "catalog of catalogs." It stores technical, business, and governance metadata for each data product (schema, owner, lineage, classification tags, quality metrics, required controls, etc.). It serves as the authoritative inventory of data products across the enterprise, enabling discovery in the data marketplace and supplying metadata to the semantic layer and control tower.
Semantic layer	The semantic layer provides shared business meaning over data products by linking data assets to common business concepts and relationships. Built on a knowledge-graph model, it connects data products in the enterprise data catalog to an enterprise ontology that defines key business terms, their definitions, and how they relate to each other.
Metadata access layer	The metadata access layer is a capability of the control tower that controls who can see data products and their metadata in the enterprise data catalog and data marketplace. It applies policy rules to determine whether a user is allowed to discover a data product at all, and which metadata fields (such as name, description, lineage, or ownership) they are permitted to view. By enforcing access controls at the metadata level, the control tower ensures that sensitive or restricted data products are not even discoverable by unauthorized users.

Component	Description
Data access layer	This is a part of the control tower that controls how approved users can access the underlying data within the data supply chains. Once a user has discovered a data product and requested access, this layer evaluates policy rules to decide whether access is allowed, denied, or subject to additional conditions such as approvals, purpose declarations, or time limits.
Data marketplace	The data marketplace provides a unified interface for discovering and accessing data products. It is a user-facing "storefront" that enables users to search and browse data products, understand their metadata (descriptions, owners, quality, policies), request access, and seamlessly consume data products. It integrates with the control tower for approval workflows and with the catalog and semantic layer to provide rich context. It does not physically store data but orchestrates controlled access and subscription to underlying data sources.

5.2 AI-ready data overarching data pattern

Aligned to the AI-ready data conceptual architecture, this pattern describes how key components, including data supply chains, the enterprise data catalog, the semantic layer, the control tower, and the data marketplace, interact cohesively to produce, govern, and consume data products within the AI-ready data platform.

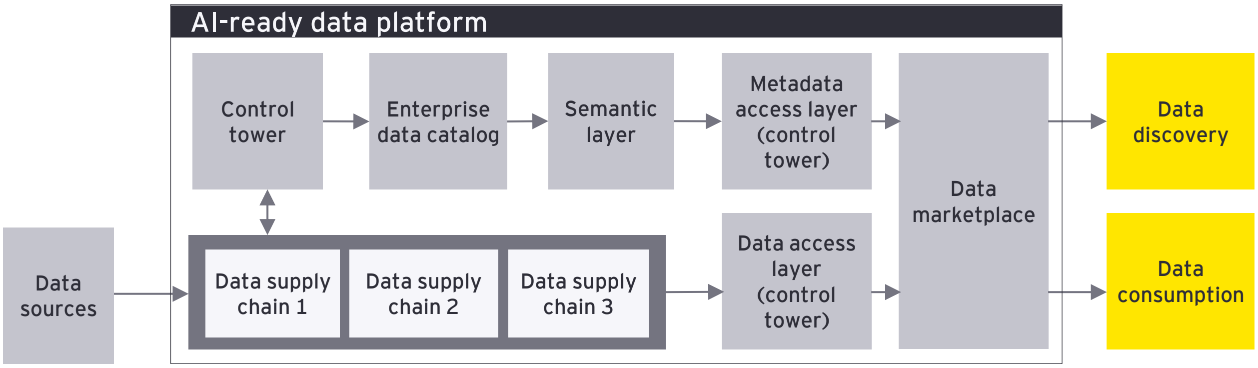

Figure 2 — AI-ready data overarching data pattern

Data is created within data supply chains, where source data is transformed into standardized, reusable data products under embedded policy controls. These supply chains remain federated (data can stay in the source or appropriate platform).

These products are registered in the enterprise data catalog, capturing metadata, lineage, and governance attributes, and are further enriched through the semantic layer to provide shared business meaning and enable cross-domain interoperability.

Governance is centrally enforced through the control tower, which applies policies, approval workflows, and control mechanisms across both data visibility and access. The metadata access layer governs what data products can be discovered based on policy and user context, while the data access layer ensures that consumption of underlying data assets adheres to defined permissions and constraints. These capabilities are surfaced through the data marketplace, which provides a unified, policy-aware interface for discovering, requesting, and consuming data products.

6. AI-ready data architecture principles

Modern AI systems demand a data foundation that is scalable, trustworthy, interoperable and capable of supporting both human-driven and autonomous intelligent workflows. These demands go well beyond traditional analytics platforms. AI-ready data architecture is built on a set of core principles that helps data to be transformed into intelligence across operational, analytical and GenAI.

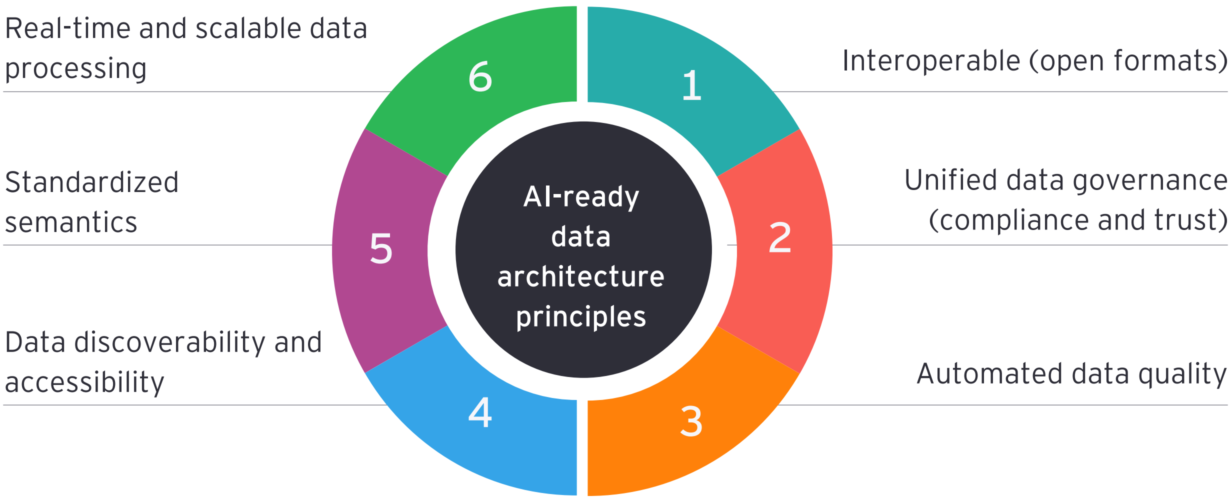

Figure 3 – AI-ready data architecture principles

Interoperable data through open formats

AI-ready data architectures avoid lock-in and fragmentation by relying on open, interoperable data formats and interfaces.

Why it matters for AI: Interoperability enables reuse, portability and faster experimentation as AI use cases evolve.

Real-time and scalable data processing

An AI-ready data foundation must react to events as they occur, processing streams, transactions, logs, sensor signals and unstructured content with low latency. Distributed, event driven architectures enable real-time insights, intelligent automation and adaptive decision-making. Scalable compute patterns support intensive workloads such as feature extraction, embeddings generation, vector search and multimodal model processing. Together, they support that the architecture grows effortlessly with business and AI demands.

Why it matters for AI: It supports real-time AI use cases and prevents AI solutions from being limited to batch analysis.

Unified data governance for compliance and trust

AI-ready architectures embed governance by design, enforcing security, privacy, lineage and life cycle controls across all data. A unified governance layer enforces policies for security, privacy, lineage, quality and lifecycle management across structured and unstructured data. This maintains that all consumers, applications, analysts and AI agents operate on trusted, well-controlled data while meeting regulatory and ethical standards.

Why it matters for AI: It reduces risk and enables confident scaling of AI into production environments.

Automated data quality for reliable AI outcomes

AI models are only as good as the data they learn from. High-quality data – accurate, consistent, complete and timely – is foundational to minimizing model bias, improving predictions, and enabling stable decision automation. Automated quality checks, anomaly detection, schema validation and continuous monitoring maintain the reliability and readiness of data feeding into model training, feature pipelines and real-time inference systems.

Why it matters for AI: It prevents silent degradation of AI performance and improves long-term model stability.

Standardized semantics for consistent interpretation

AI-ready data foundation requires more than raw data; it requires shared meaning. A unified semantic layer standardizes business definitions, metrics, hierarchies and entity relationships across the enterprise. This eliminates inconsistencies between analytics, operational systems and AI models, enabling that both humans and autonomous agents interpret data correctly. Semantic consistency accelerates feature engineering, improves model explainability, and embeds domain logic into AI pipelines.

Why it matters for AI: It improves explainability, consistency and reuse across analytics and AI systems.

Data discoverability and accessibility

AI-driven innovation depends on the ability to quickly find and access relevant data sets. A strong metadata layer, a searchable catalog and governed access paths make data discoverable to both humans and AI agents. Standardized APIs, semantic endpoints and federated access patterns reduce friction, enabling rapid experimentation, reuse and cross-domain intelligence. This principle supports the right data is always within reach, without compromising security.

Why it matters for AI: It accelerates innovation while maintaining security and compliance.

Together, these principles define an architecture designed for AI from the outset, rather than adapted from legacy analytics platforms.

7. Modern data supply chain capabilities for enterprise AI

Modern AI data supply chain capabilities represent a fundamental shift from traditional data platforms to an intelligent, always-on ecosystem that powers enterprise AI at scale. Rather than operating as fragmented tools or pipelines, these capabilities form a cohesive, end-to-end architecture that enables continuous ingestion, processing, governance, and delivery of high-quality, context-rich data in real time. This integrated data foundation is essential for supporting advanced AI workloads, from analytics and machine learning to agentic systems, where speed, scalability, interoperability, and trust are critical.

At its core, this architecture comprises interoperable building blocks that transform raw, fragmented data into reliable, AI-ready assets. Spanning multi-source ingestion, real-time intelligent processing, unified storage, governed access, semantic consistency, and self-service consumption, each layer ensures seamless data flow across the enterprise. Collectively, these capabilities accelerate AI innovation, reduce time-to-insight, and enable confident operationalization of AI—driving consistent, scalable, and business-aligned outcomes.

7.1 Data ingestion and integration

This building block maintains that all relevant data – structured, semi structured, unstructured and streaming – can be ingested, integrated and made available for analytics.

Core capabilities:

- Multisource data ingestion (batch, streaming, change data capture) source data ingestion (batch, streaming, change data capture)
- Integration from APIs, IoT, SaaS, on-prem, cloud systems
- Schema evolution support
- Real-time and micro-batch data processing

Why it matters for AI

AI workloads demand high-volume and high-variety data sets; this building block enables seamless, timely access to all required data types.

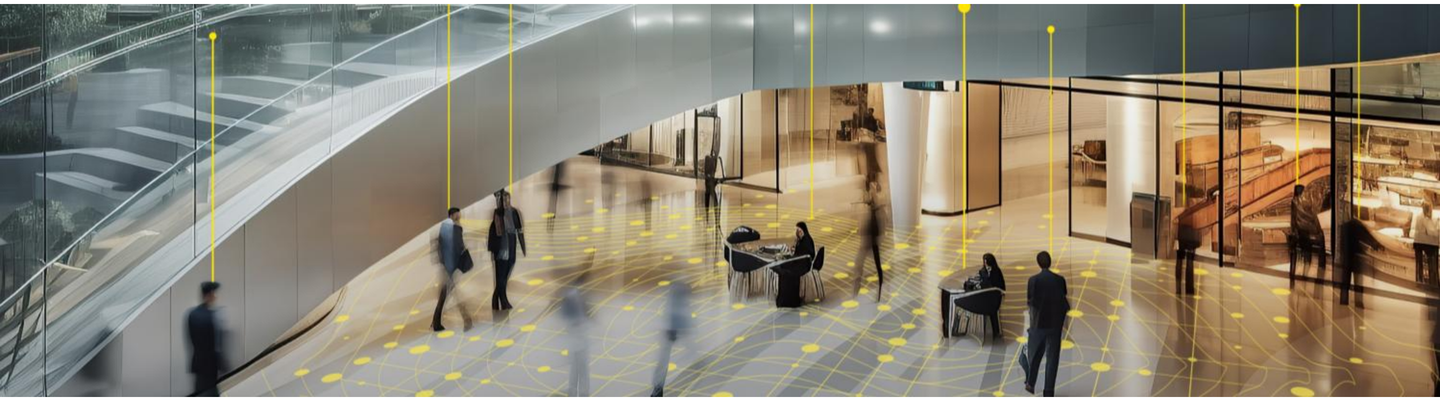

Data ingestion patterns

Below are the various data ingestion patterns:

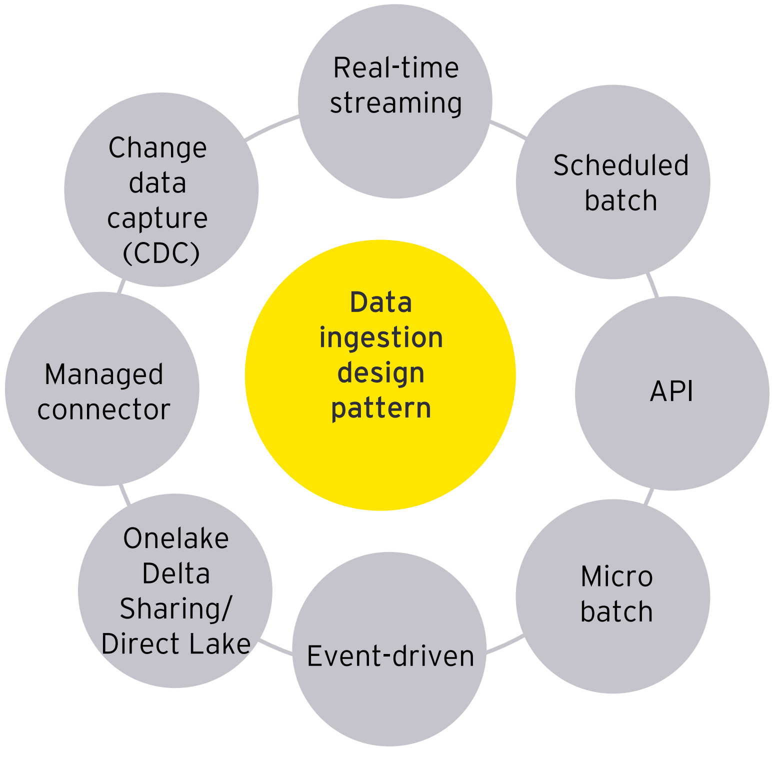

Figure 4 — Data ingestion patterns

7.2 Intelligent, real-time and scalable data processing

Modern AI systems require the ability to process structured and unstructured data continuously, intelligently and at scale. This building block integrates real-time ingestion, distributed processing and advanced feature engineering to turn raw data – from tables to text, images, logs and audio – into AI-ready assets suitable for training and real-time inference, rather than static data sets.

Core capabilities:

- **Scalable processing for unstructured data**
  This capability leverages distributed and vectorized processing engines to enable embedding generation, semantic search and feature extraction, unlocking the full value of NLP, computer vision and multimodal workloads, without being tied to a specific vendor or platform.
- **Distributed data processing engines**
  Spark, SQL, and Python pipelines support large-scale ETL or ELT, accelerating the preparation of data sets for both model training and real-time inference.
- **Reusable, reproducible feature engineering**
  Managed workflows support consistency across training and production, enabling versioned, reliable and shareable features that can be reused across teams and AI use cases.

Why it matters for AI

Feature engineering often represents 70%+ of ML effort, making intelligent, scalable data processing essential for shortening experimentation cycles, improving model quality, and reducing friction when moving AI solutions from development to production.

Data processing patterns and activities:

Medallion layer architecture

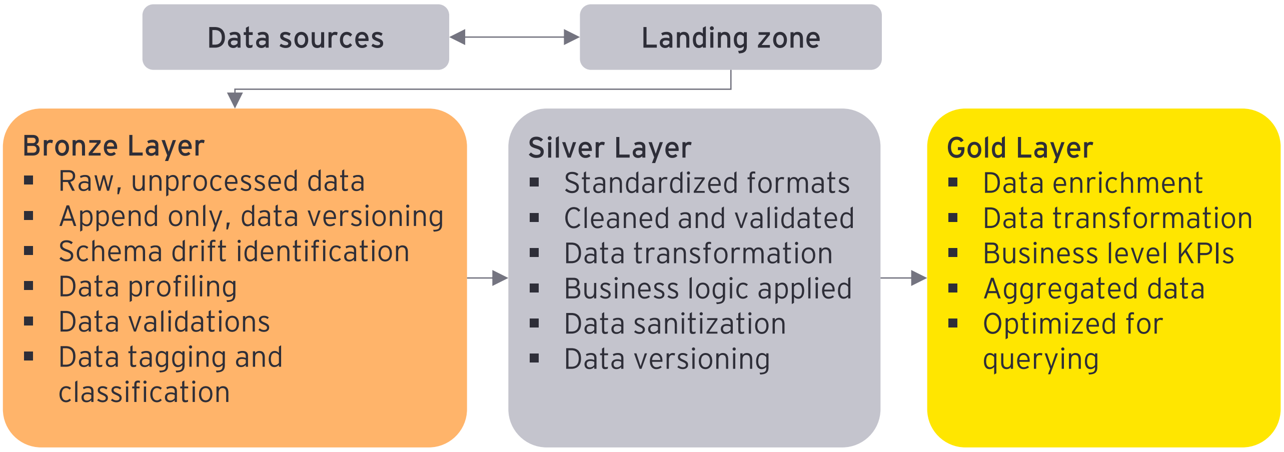

Figure 5 – Data processing patterns and activities

In many enterprise architectures, raw data is first ingested into a controlled processing zone where security, quality and compliance checks are applied before downstream transformation. An initial raw data layer typically preserves immutable source data to maintain lineage, enable replay and support audit and traceability requirements. A subsequent standardized layer applies validation, cleansing and transformation logic to produce trusted, analytics- and AI-ready data sets. A curated consumption layer then delivers enriched and ready-to-use data sets for analytics, AI features, and operational decision-making.

7.3 Unified data and feature foundation

This is a modern unified foundation that brings together Lakehouse-style storage and feature management to support analytics and AI at scale. It combines the flexibility of data lakes, the reliability of data warehouses, and the operational consistency of feature stores into a single architectural foundation rather than as separate platforms.

Core capabilities:

- Single logical data layer for all formats: it supports structured, semi structured (JSON, Avro, Parquet) and unstructured data (images, video, audio, logs, documents), allowing AI workloads to access diverse data types without fragmentation.
- Reliable and well-governed storage: ACID transactions, indexing and advanced Lakehouse formats such as Delta and Iceberg provide consistency, performance and efficient storage on cost-efficient cloud object storage.
- Separation of storage and compute: it enables elastic scaling and cost-efficient performance by allowing compute resources to scale independently of data storage.
- Feature store integration: it provides versioned, reusable feature pipelines with real-time and batch serving, maintaining consistent features across training and inference and reducing discrepancies between development and production models.

Why it matters for AI

AI workloads require massive historical data sets, multiple data set versions and native support for diverse unstructured content. The unified lakehouse + feature store architecture delivers the scale, structure and consistency needed for faster experimentation, reliable training and consistent model behaviour across the full AI lifecycle.

Data processing patterns and activities:

Data store patterns

AI-ready data architectures typically employ multiple data store patterns, selected based on workload characteristics, performance needs and access requirements. Common patterns include:

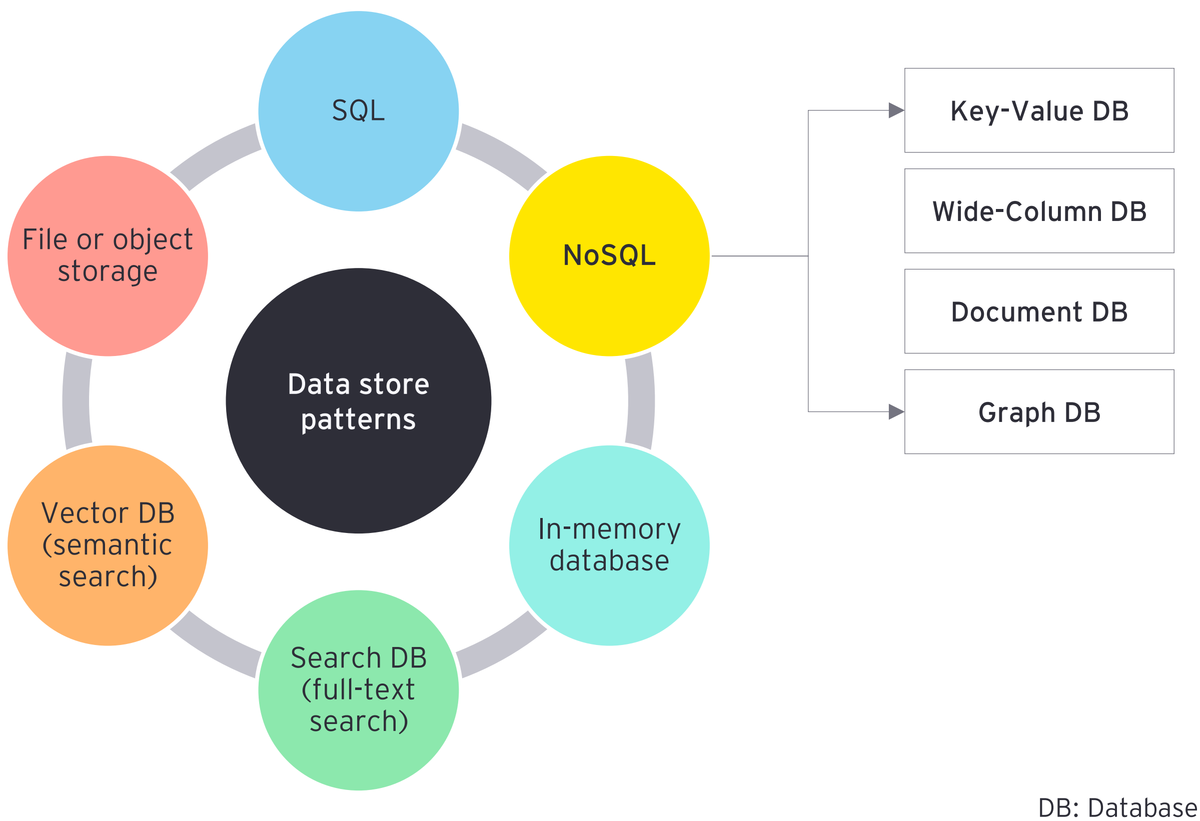

Figure 6 – Data store patterns

Figure 6 highlights the primary data storage patterns commonly used in AI-ready architectures described below:

1. SQL databases, for relational and transactional workloads
2. NoSQL databases, including:
    - Key-Value database
    - Wide-Column database
    - Document database
    - Graph database
3. File or object storage for large-scale structured and unstructured data
4. Vector database (semantic search) to support semantic search and retrieval-augmented AI use cases
5. Search database (full-text search) for full-text indexing and retrieval
6. In-memory database for low-latency and high-throughput access patterns

7.4 Data quality, governance and security

This building block establishes trust, control and compliance across the data lifecycle by embedding governance, quality and security consistently across all data flows and consumption patterns.

Core capabilities:

- Automated data quality validation and anomaly detection to identify data issues early and prevent defects from propagating into AI models
- Data catalog, lineage tracking, business glossary to provide transparency into data origins, transformations and meaning
- Role-based access control (RBAC), Attribute-based access control (ABAC), and Policy enforcement-based access control (PBAC) to enforce fine-grained and scalable authorization across diverse users, applications and AI systems
- Data masking, tokenization, encryption to protect sensitive data across storage, processing and consumption
- Compliance support (GDPR, HIPAA, regional regulations) to maintain AI workloads operating within regulatory and data residency requirements

Why it matters for AI

Without trusted data, models degrade quickly, predictions become unreliable, and regulatory exposure increases. As AI adoption scales, weaknesses in data quality, governance or security are amplified, directly impacting model performance, explainability and organizational confidence in AI outcomes.

Data authorization patterns

AI-ready data architectures typically apply multiple authorization patterns in combination, balancing simplicity, flexibility and scale. Common patterns include:

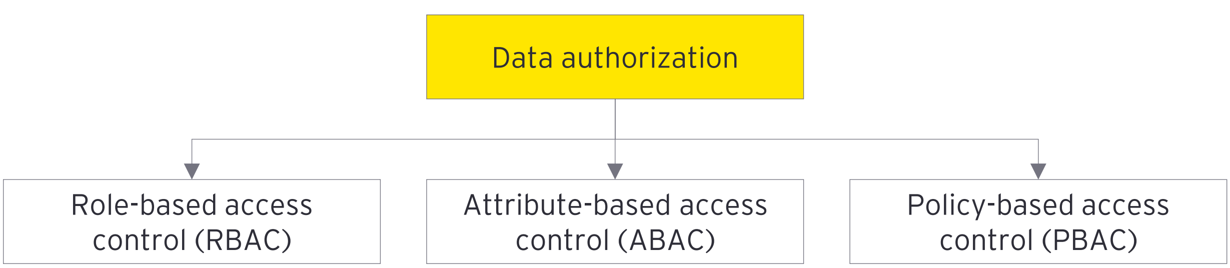

Figure 7 – Data authorization design patterns

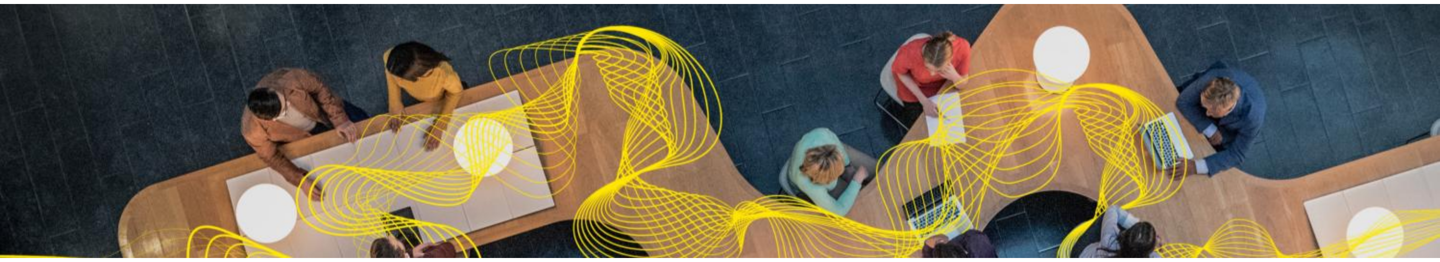

Below are the three data authorization design patterns explained in figure 7:

- **Role-based access control (RBAC):** RBAC assigns permissions based on predefined organizational roles. Users inherit access by being mapped to roles like analyst or admin. It simplifies management, supports standard job functions, and works well in stable, hierarchical environments. However, it becomes rigid when granular or dynamic access rules are needed.
- **Attribute-based access control (ABAC):** ABAC makes access decisions using attributes about users, data, actions and context. Policies evaluate factors like department, sensitivity, time or location. It offers fine-grained, flexible, dynamic control suited for complex authorization needs but requires strong governance and consistent attribute metadata.
- **Policy-based access control (PBAC):** PBAC centralizes authorization through policies evaluated by a policy engine. Policies define who can access what and under what conditions, combining roles, attributes and business rules. It enables consistent, scalable, cross-system governance, though it relies on mature policy management and clear rule ownership.

7.5 Standardized semantics with a unified semantic layer

A semantic layer supports that all analytical, operational, and AI workloads interpret data consistently, regardless of the tools, teams or downstream systems accessing it. By providing standardized business definitions, metrics and relationships, it becomes the translation layer between raw data and meaningful, governed insights so that AI and analytics operate on the same business meaning.

Core capabilities:

- Centralized business definitions and metrics: They support that KPIs, dimensions, entities and hierarchies are defined once and reused everywhere, across dashboards, SQL queries, ML pipelines and APIs to reduce duplication and inconsistencies across teams.
- Semantic modelling for structured, semi structured and unstructured data: It attaches meaning, metadata and relationships to diverse data sets, enabling AI systems to understand entities, context and domain logic in a consistent and governed way.
- Consistent data consumption across tools: BI tools (e.g., Power BI), ML platforms and custom apps use the same semantic definitions, eliminating mismatches between training data and production logic and improving consistency between experimentation and real-world deployment.
- Governed access and policy enforcement: Row-level security, data masking and glossary-driven governance maintain responsible, compliant usage of enterprise data while maintaining a consistent interpretation of business rules.

Why it matters for AI

AI models depend on consistent, context-rich and well-defined data. A semantic layer provides:

- Clear lineage and meaning behind features
- Standardized entity relationships needed for feature engineering
- Reduced data ambiguity across teams
- Easier integration of AI with business logic

This building block elevates raw data into trusted, context-aware AI signals, accelerating model development and supporting alignment with business outcomes as AI scales across domains and teams.

7.6 Self-service data democratization for applications, users and AI agents

This building block focuses on enabling seamless, governed and autonomous access to enterprise data for a wide range of consumers, including business applications, human users and increasingly AI agents. By democratizing both metadata and actual data access, it enables frictionless consumption across analytical, operational and AI-driven workflows while maintaining appropriate governance and control.

Core capabilities:

- **Metadata-driven data discovery:** A centralized data marketplace and rich metadata – schema, lineage, data quality and business definitions – enable users and AI agents to efficiently discover, evaluate, and request data sets with full governance and transparency, so consumers can find and use trusted data without manual coordination.
- **Multi-modal data consumption:** Trusted data is consumed through various patterns, including business applications, self-service analytics, semantic APIs, dashboards and agentic AI systems that autonomously retrieve data, generate insights, and act on decisions, using consistent definitions and governed access paths.
- **Unified access via data virtualization:** A logical access layer provides a single, real-time view of distributed data without replication. It enables secure, cross-domain analytics through SQL, APIs and semantic interfaces while minimizing movement and improving agility where data movement is costly or constrained.

Why it matters for AI

As AI agents become first-class consumers of enterprise data, democratized access provides availability, usability and trust. It enables both humans and AI systems to retrieve the right data at the right time, accelerating continuous intelligence, automation and business impact across the organization without increasing governance risk.

Data consumption design patterns:

In many enterprises, data consumption is typically organized into two complementary patterns: metadata consumption (discovery) and actual data consumption (use), delivered through multiple governed connectivity methods.

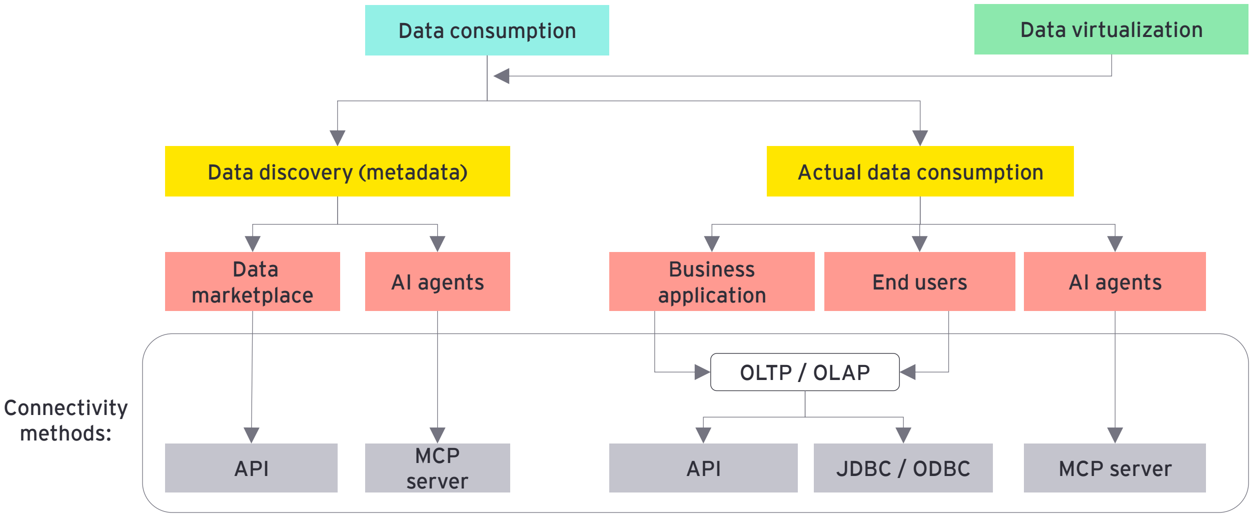

Figure 8 – Data consumption design patterns

As illustrated in figure 8, AI-ready data architectures support multiple data consumption patterns, as explained below:

- **Data virtualization:** It provides a logical data access layer that enables consumers to query and retrieve data from multiple sources without physically moving or replicating it. This pattern is essential for scenarios requiring real-time access, cross-domain analytics and minimal data movement. By abstracting underlying complexities, data virtualization supports faster delivery, reduced storage costs and improved agility while maintaining governance and security. It supports consumption through standardized interfaces like SQL, APIs and semantic layers, making it ideal for business applications, AI agents and self-service analytics.
- **Metadata consumption (data discovery):** Metadata consumption is typically enabled through a data catalog or a marketplace experience and can also support AI-assisted discovery where appropriate.
- **Actual data consumption:** Actual data consumption involves accessing and using the underlying data for business purposes. This includes:
  - **Business applications:** OLTP or OLAP application through API or JDBC or ODBC connectivity
  - **Agentic AI solutions:** AI agents accessing enterprise data through governed agent-to-data connectors (e.g., an agent connectivity server pattern)
  - **End users:** analysts, data scientists and business teams leveraging dashboards, self-service analytics through APIs or SQL based connectivity methods
- **API connectivity:** RESTful or GraphQL APIs provide standardized, scalable and secure access to data services. APIs are widely used for business applications, AI and agentic workflows, and integration with external systems, providing real-time or near-real-time data delivery.
- **JDBC or ODBC connectivity:** These traditional connectivity protocols allow direct access to relational and analytical data stores from BI tools, reporting platforms and custom applications. These are commonly used by end-users for self-service analytics and by business applications for structured data queries.
- **Agentic connectivity server pattern:** A governed connectivity server can provide secure, high-performance connections between AI agents and enterprise data services, supporting compliance requirements and reliable throughput for high-volume use cases.

Bridge: from architecture building blocks to business value

Together, these building blocks translate AI-ready data architecture principles into an operational foundation that enables real business outcomes. By combining interoperable ingestion, scalable processing, unified data and feature foundations, embedded governance, standardized semantics and governed self-service access, enterprises reduce friction across the AI lifecycle and create the conditions for AI to scale safely. This architectural coherence is what underpins the benefits described earlier — faster time-to-value, higher productivity, improved trust and the ability to deploy AI confidently across critical workflows. In practice, the organizations that realize sustained AI value are not those with the most advanced models, but those with data architectures deliberately designed to support speed, scale and trust simultaneously.

8. Business value and outcomes

An AI-ready data architecture enables high quality, governed, interoperable and real time data, enabling organizations to fully realize the value of AI. AI-ready data foundations directly shape whether AI initiatives scale beyond pilots and deliver sustained business impact.

Faster time-to-value and higher return on AI investments

Traditional data architecture introduces friction through duplicated pipelines, manual data preparation and disconnected governance controls. AI-ready foundations replace this with reusable components, shared semantics and standardized access patterns.

Business impact: AI initiatives reach production faster, experimentation cycles shrink, and organizations extract value from AI investments sooner and more predictably.

Increased productivity across data and AI teams

AI-ready architectures automate many labor-intensive activities such as data integration, quality validation, feature preparation and access provisioning. This reduces the operational burden on engineering and data science teams.

Business impact: Teams spend more time building and improving AI solutions and less time on data preparation and fixing data issues, resulting in faster onboarding of new data sources, measurable gains in delivery velocity and efficiency, reduced time-to-insight for analytics and GenAI workloads.

Improved trust, transparency and risk management

As AI adoption grows, trust becomes a decisive factor. AI-ready data platforms embed lineage, quality signals, governance and access controls across all data flows, including unstructured and real-time data.

Business impact: Leaders gain confidence in AI outputs, regulators can be satisfied more easily, and organizations reduce the risk of deploying models that are opaque, non-compliant or unreliable.

Consistent business meaning across AI and analytics

Without standardized semantics, AI systems and analytics often operate on different interpretations of the same data. AI-ready architectures establish a shared semantic layer that aligns data definitions, metrics and relationships across tools and use cases.

Business impact: AI outputs are easier to explain, validate and act upon, improving adoption and decision-making across business functions.

Enablement of advanced and real-time AI use cases

AI-ready data foundations support real-time pipelines, multimodal data and scalable processing needed for modern AI, including GenAI, retrieval-augmented generation and autonomous agents.

Business impact: Organizations can move beyond offline analysis to deploy AI that informs and automates decisions in operational workflows.

Long-term scalability and lower cost of change

AI demand grows unpredictably. Architectures built specifically for AI scale more efficiently over time by avoiding repeated re-engineering as new models, data sources and regulations emerge.

Business impact: Technology spend becomes more predictable, platforms remain adaptable, and AI capabilities can expand without destabilizing the core data environment.

These benefits explain why AI-ready data architecture is increasingly seen as a prerequisite for scaling AI responsibly, enabling speed and innovation without sacrificing trust, compliance or control.

9. Implementation roadmap for AI-ready transformation

An AI-ready data architecture is no longer optional, it's a foundational enabler for scalable, trusted and responsible AI as organizations move from experimentation to enterprise-wide adoption. Organizations can begin unlocking this value through a structured, assessment-led engagement approach grounded in architecture, governance and business priorities.

Suggested first steps

To rapidly progress toward an AI-ready data architecture, organizations should begin with the following sequenced steps, focusing on learning, validation and incremental scale rather than a single large transformation:

Phase 1. Conduct a data architecture maturity assessment:

Perform an AI-ready data architecture maturity evaluation across governance, semantics, data engineering, interoperability, unstructured data readiness and compliance to establish a clear baseline and identify priority gaps.

A structured, end-to-end AI-ready data architecture assessment can be used to evaluate:

- The current data estate, including structured, unstructured, multimodal data readiness
- Data governance, quality, lineage and metadata maturity across analytics and AI workflows
- Semantic layer, ontology and interoperability capabilities supporting consistent interpretation
- Real-time data availability and engineering patterns required for low-latency AI use cases
- Alignment with business use cases and AI adoption roadmap
- Gaps between traditional architecture and AI-ready operating models

This assessment establishes a clear baseline and provides a prioritized, actionable roadmap for data modernization and AI enablement, helping organizations focus investment where it will have the greatest impact.

Phase 2. Identify priority AI use cases:

Select high-value business use cases that require trusted, timely AI-ready data as foundational enablers (e.g., GenAI copilots, RAG, forecasting, customer analytics) to anchor architectural decisions in business outcomes.

Phase 3. Launch a pilot workstream:

Initiate a targeted pilot (e.g., unstructured data onboarding, semantic layer build-out, vector-ready data pipeline or domain data product) to validate architecture patterns and demonstrate measurable value early.

Phase 4. Build the AI-ready data roadmap:

Define a phased modernization path that includes standards, governance models, platform enhancements, semantic strategy and operating model changes to guide consistent adoption across domains and teams.

Phase 5. Scale and industrialize:

Scale proven patterns using standardized reference architectures, templates and automation approaches to industrialize AI-ready foundations across domains and business units while maintaining consistency, governance and operational efficiency.

10. Conclusion

An AI-ready data architecture is now a strategic necessity for organizations aiming to scale AI with trust, agility and measurable business value. This whitepaper outlined the principles, architectural components and leading practices required to modernize enterprise data ecosystems into resilient, intelligent and future proof foundations that support advanced analytics, machine learning and emerging AI agent workloads.

As models grow in complexity and data volumes span structured to unstructured domains, organizations must move beyond traditional data management. Modern AI demands real-time data pipelines, open and interoperable formats, embedded governance, semantic enrichment, automated data quality and broad self-service access, increasingly extending even to autonomous AI agents.

Building such an architecture is a continuous, incremental journey requiring tight alignment across technical, governance and business teams, strong engineering discipline and commitment to standards. Organizations adopting these principles can operationalize AI responsibly, move from experimentation to production with confidence and unlock long-term, compounding value from their data assets.

Key considerations for implementing AI-ready data architecture

Implementation complexity and organizational change management:
Modernizing into an AI-ready architecture requires significant process transformation, alignment across teams and cultural adaptation. Without structured change management, adoption becomes fragmented and slows progress.

Skills gaps and talent requirements:
AI-ready environments rely on capabilities spanning data engineering, MLOps, governance and cloud architecture. A shortage of these skills can impede successful implementation and limit value realization.

Cost considerations for migrating from legacy architectures:
Transitioning from legacy, siloed systems to modern, scalable platforms introduces substantial costs in tooling, infrastructure, data migration and workforce enablement, often underestimated in early planning.

Vendor lock-in risks despite open format adoption:
Even with open data formats, platform services, proprietary optimizations and cloud-native integrations can create long-term dependencies that constrain flexibility and portability.

Data quality debt prior to AI readiness:
Unresolved data quality issues, such as inconsistencies, missing metadata or lack of lineage, must be addressed early. Accumulated data debt undermines model accuracy, governance and trust in AI outcomes.

Organizations that treat data architecture as a strategic foundation, not a technical afterthought, will be the ones that successfully scale AI and sustain competitive advantage.

Authors

Graham Gibson

Associate Partner, Data Platform Leader, Global Data Office, EY Growth & Innovation, Sydney, Australia
Ernst & Young Services Pty Limited

Ramana Reddy Burri

Assistant Director, EY Technology, Client Technology Engineering — Data Competency
Bangalore, India
EY Global Delivery Services (India) LLP