tieto

From data readiness to trusted AI

Building scalable, enterprise-ready AI in six steps

Visit tieto.com

Table of contents

Introduction

01. Start with value: align data to AI use cases

Problem 6
Recommended approach 6
Case study: Unlocking production intelligence through unified data 7
Executive takeaway 7

02. Building an AI-ready data platform

Problem 9
Recommended approach 9
Case study: Cloud-native data foundation for AI 10
Executive takeaway 10

03. Embedded data governance as a foundation for trusted AI

Problem 12
Recommended approach 12
Executive takeaway 13

04. Data preparation for AI

Problem 15
Recommended approach 15
Case study: A phased approach for supervised learning 16
Executive takeaway 16

05. Model training and deployment: balancing automation with human oversight

Problem 18
Recommended approach 18
Case study: Sustaining AI performance through MLOps 19
Executive takeaway 19

06. Ensuring AI adoption throughout the organization

Problem 21
Recommended approach 21
Case study: Securing AI adoption at scale 22
Executive takeaway 22

Conclusion and outlook: data readiness in the age of agentic AI

About Tieto Tech Consulting

Introduction

As organizations accelerate their adoption of artificial intelligence (AI), data has become one of the most important strategic assets. Yet, despite growing investments in AI technologies, many initiatives still fail to deliver the expected value, and in most cases, the issue is not the model itself but the readiness of the underlying data. As per a report by McKinsey, eight in ten companies quoted data limitations as a roadblock to scaling agentic AI to deliver real value.

Data readiness describes the extent to which data is prepared, governed, and fit for use. It is a multidimensional concept that encompasses technical, operational, legal, and ethical considerations. At its core, data readiness depends on six essential dimensions:

01

Data quality

Accuracy, completeness, consistency, and timeliness of data

02

Accessibility

Ability to retrieve and use data efficiently across systems and teams

03

Governance

Clear policies around data ownership, stewardship, and usage rights

04

Metadata and lineage

Transparency and documentation of data origin, transformations, and context

05

Security and privacy

Protection of sensitive data and compliance with regulatory frameworks such as GDPR or the EU AI Act

06

Ethical integrity

Bias identification and mitigation, fairness, and responsible use.

Together, these dimensions determine whether data is available, trustworthy, compliant, and contextually rich enough to support decisions.

Organizations that neglect them face more than technical setbacks. Poor data readiness can lead to inaccurate outputs, slow delivery, rising costs, limited user trust, regulatory exposure, and missed business opportunities. By contrast, companies that invest in data readiness create conditions for faster time to value, stronger user adoption, lower delivery risk, and more reliable AI outcomes.

To unlock the full potential of AI, data readiness must therefore be treated as a foundational business capability rather than a technical afterthought. This development can be structured into six practical steps. In this white paper, we explore each of them and illustrate how organizations can build a data foundation that is truly ready for AI.

Six steps to scalable and trusted AI

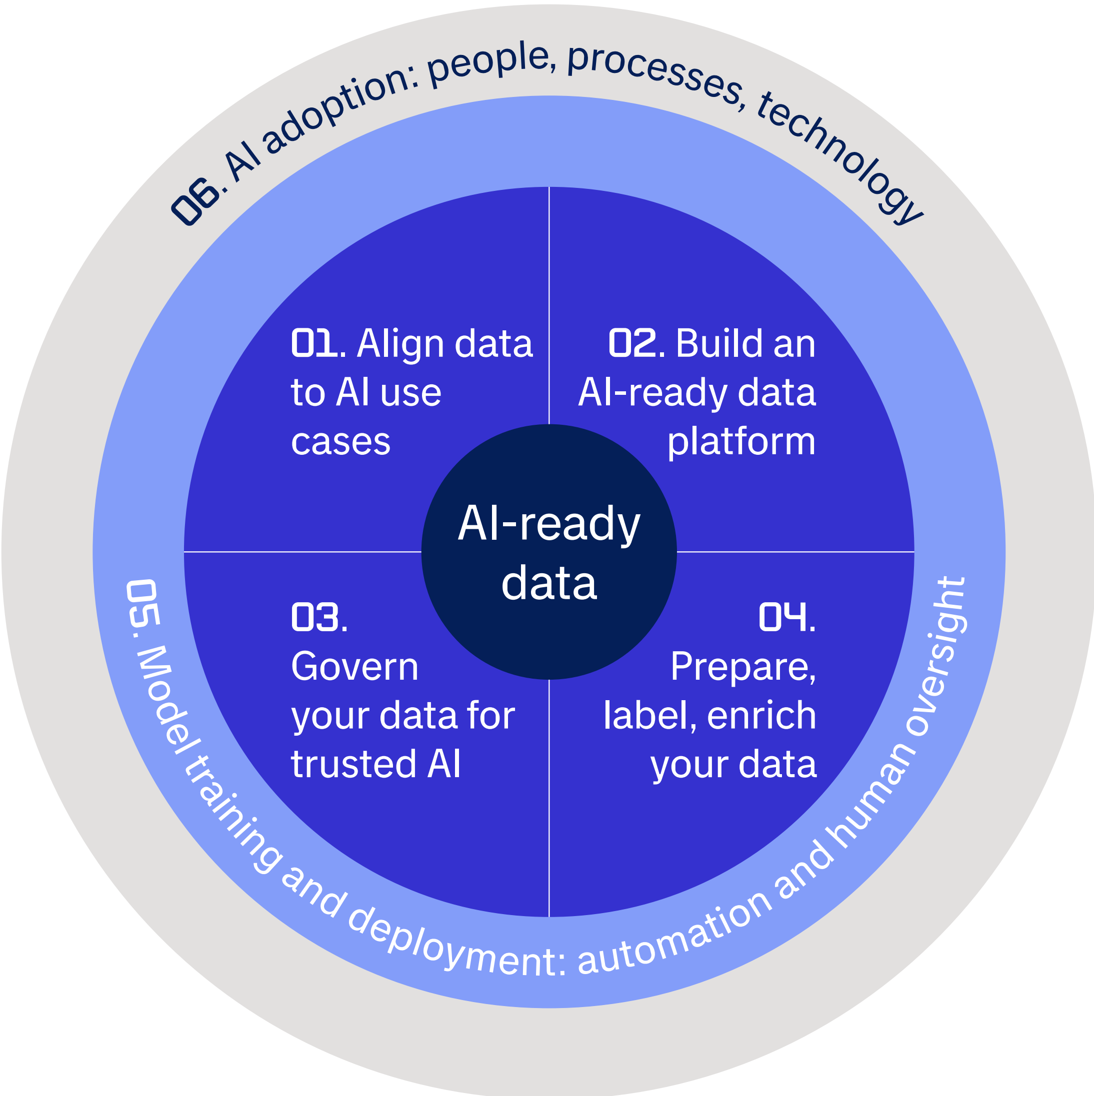

01

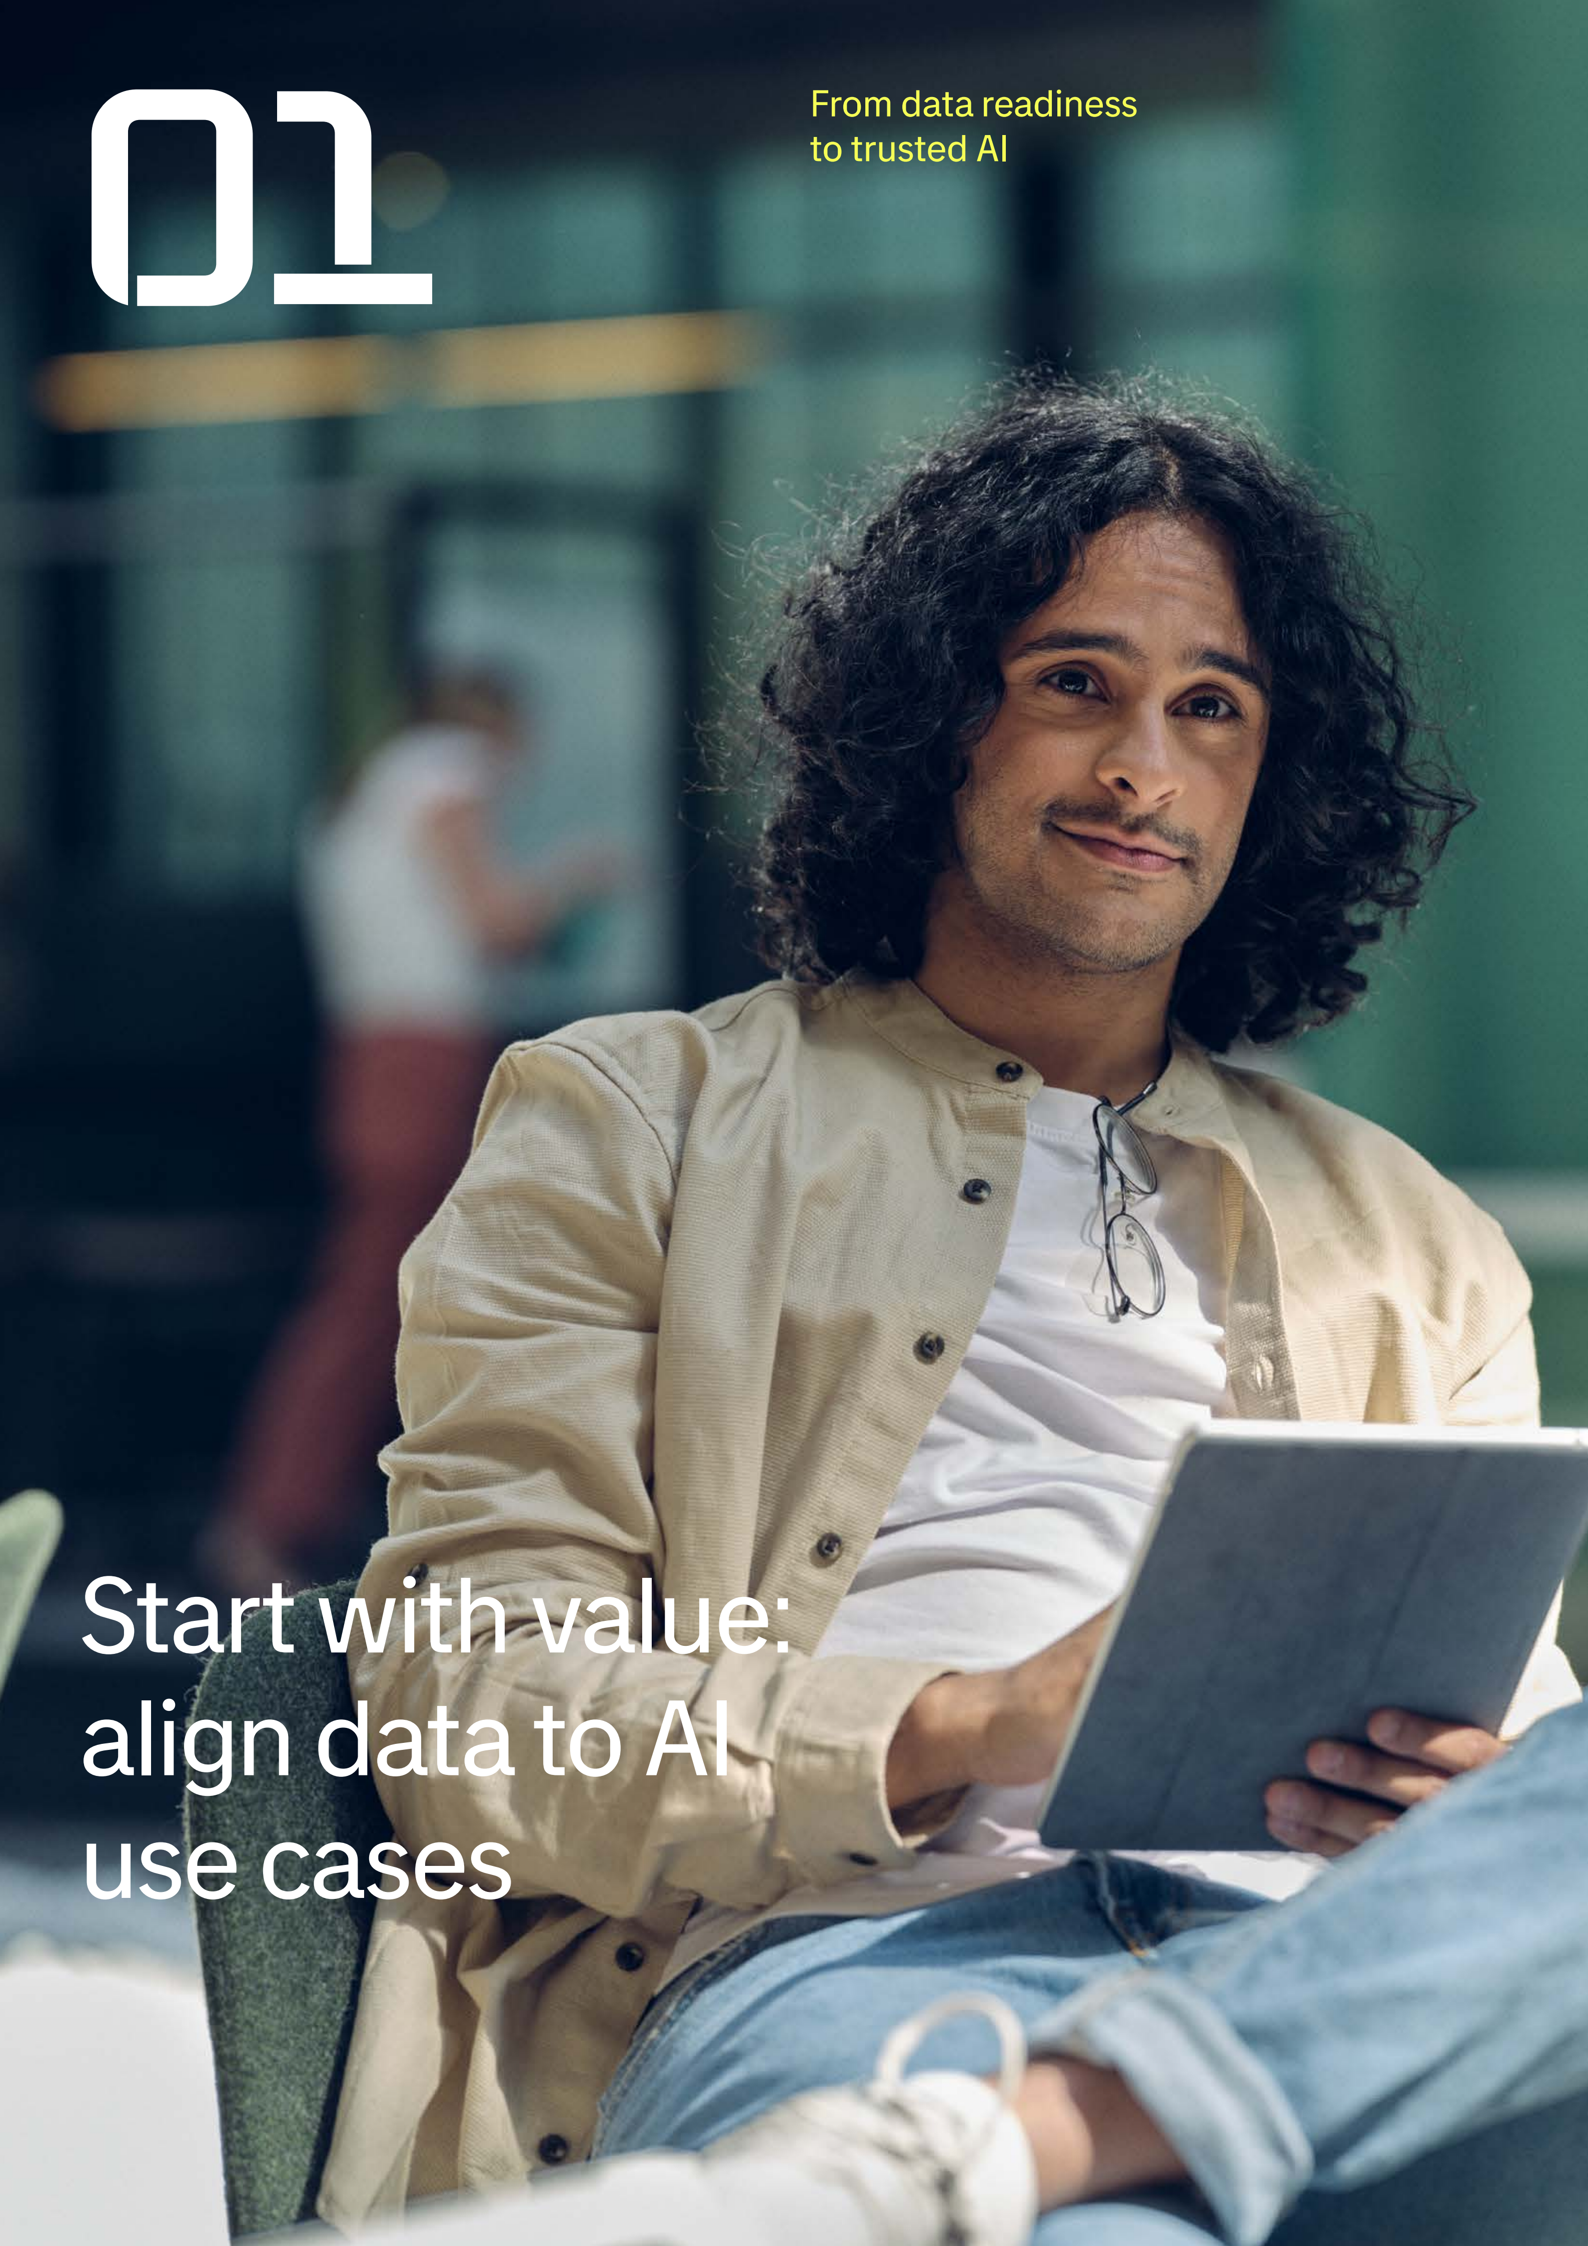

Start with value: align data to AI use cases

Problem

Many organizations approach AI by starting with technology rather than outcomes. They invest heavily in tools, platforms, and models but struggle to generate measurable business impact. The root cause is consistent across industries: AI initiatives are launched without a clear link between the business problem, the required data, and the intended value.

Three issues repeatedly undermine AI value realization:

- Use cases are defined at a conceptual level, but the underlying data generation process is not understood early enough, leading to incomplete, inaccessible, or missing data.
- Existing data is created for operations or reporting and often lacks granularity, structure, or context to be usable by AI.
- Business stakeholders assume the available data is "good enough", but AI requires significantly higher levels of trust, clarity, and completeness than traditional analytics to rely on the proper business context.

As a result, many AI projects stall in the earliest phases because business strategy, data foundations, and operational realities are not aligned.

Recommended approach

Shifting from technology-driven to value-driven AI requires business stakeholders to anchor every initiative to clear business outcomes and the data required to achieve them. The fastest way to get to successful AI is to start with clarifying which business process needs to be improved and what measurable outcome must be achieved, then work backwards to the data required.

Organizations should therefore:

- Define business value and identify which decisions or process steps must be improved with AI as well as the expected value creation. Here it is important to establish Key Performance Indicators (KPIs) to measure what matters most and make explicit "what good looks like" in terms of expected AI output, including accuracy, timeliness, and risk.
- Translate the use case into clear data requirements by involving subject matter experts to understand how decisions are made in practice, where the corresponding data is generated, and how much of it sits in silos or informal workarounds.
- Map the relevant sources and assess whether the data can be accessed responsibly and used consistently.
- Identify missing, siloed, or low-quality data early and reassess the feasibility of the use case where necessary.

Not every AI use case is feasible with the current data, and business stakeholders must be able to make strategic decisions to close the gaps (invest in data gathering, foster data improvement, or re-scope the AI use case), so that AI results are based and delivered on evidence, not assumptions or hallucinations.

Case study

Unlocking production intelligence through unified data

Tieto Tech Consulting enabled a leading European construction company to accelerate its digital transformation through a modern, standard-based data management solution. The company operates more than 200 production sites with over 1,000 production lines across Europe. While data was already available and sufficient for reporting, it quickly became clear that it lacked the granularity required for advanced analysis and modelling.

The level of digital maturity varied significantly from plant to plant. Some sites were already advanced, while others captured only minimal information. This imbalance meant that, although management reports could be produced, the data was too aggregated and inconsistent to support more ambitious goals such as process optimization, energy efficiency analysis, or predictive use cases. As soon as the business wanted to move beyond descriptive reporting and gain deeper operational insights, additional and more granular data became essential.

Our team helped establish a unified and scalable data foundation across all sites and production lines that consistently collected machine and process data and made it available centrally. With this foundation in place, the company gained transparent, comparable insights across sites and was able to move from basic reporting to data-driven decision making such as production optimization or reduction of energy consumption.

Executive takeaway

AI creates value only when business priorities, data reality, and governance are aligned from the start. Organizations that begin with a clear business objective and assess data feasibility early are better positioned to scale AI in a way that is credible, practical, and outcome-driven.

Before making any significant investments, business and IT stakeholders must provide honest answers to these questions to ensure the conditions for success exist:

What concrete business decision or process will this AI use case improve?

Is AI truly required, or would better reporting, automation, or rules-based logic already be able to solve the problem?

Are there critical data gaps that would fundamentally undermine the initiative?

02

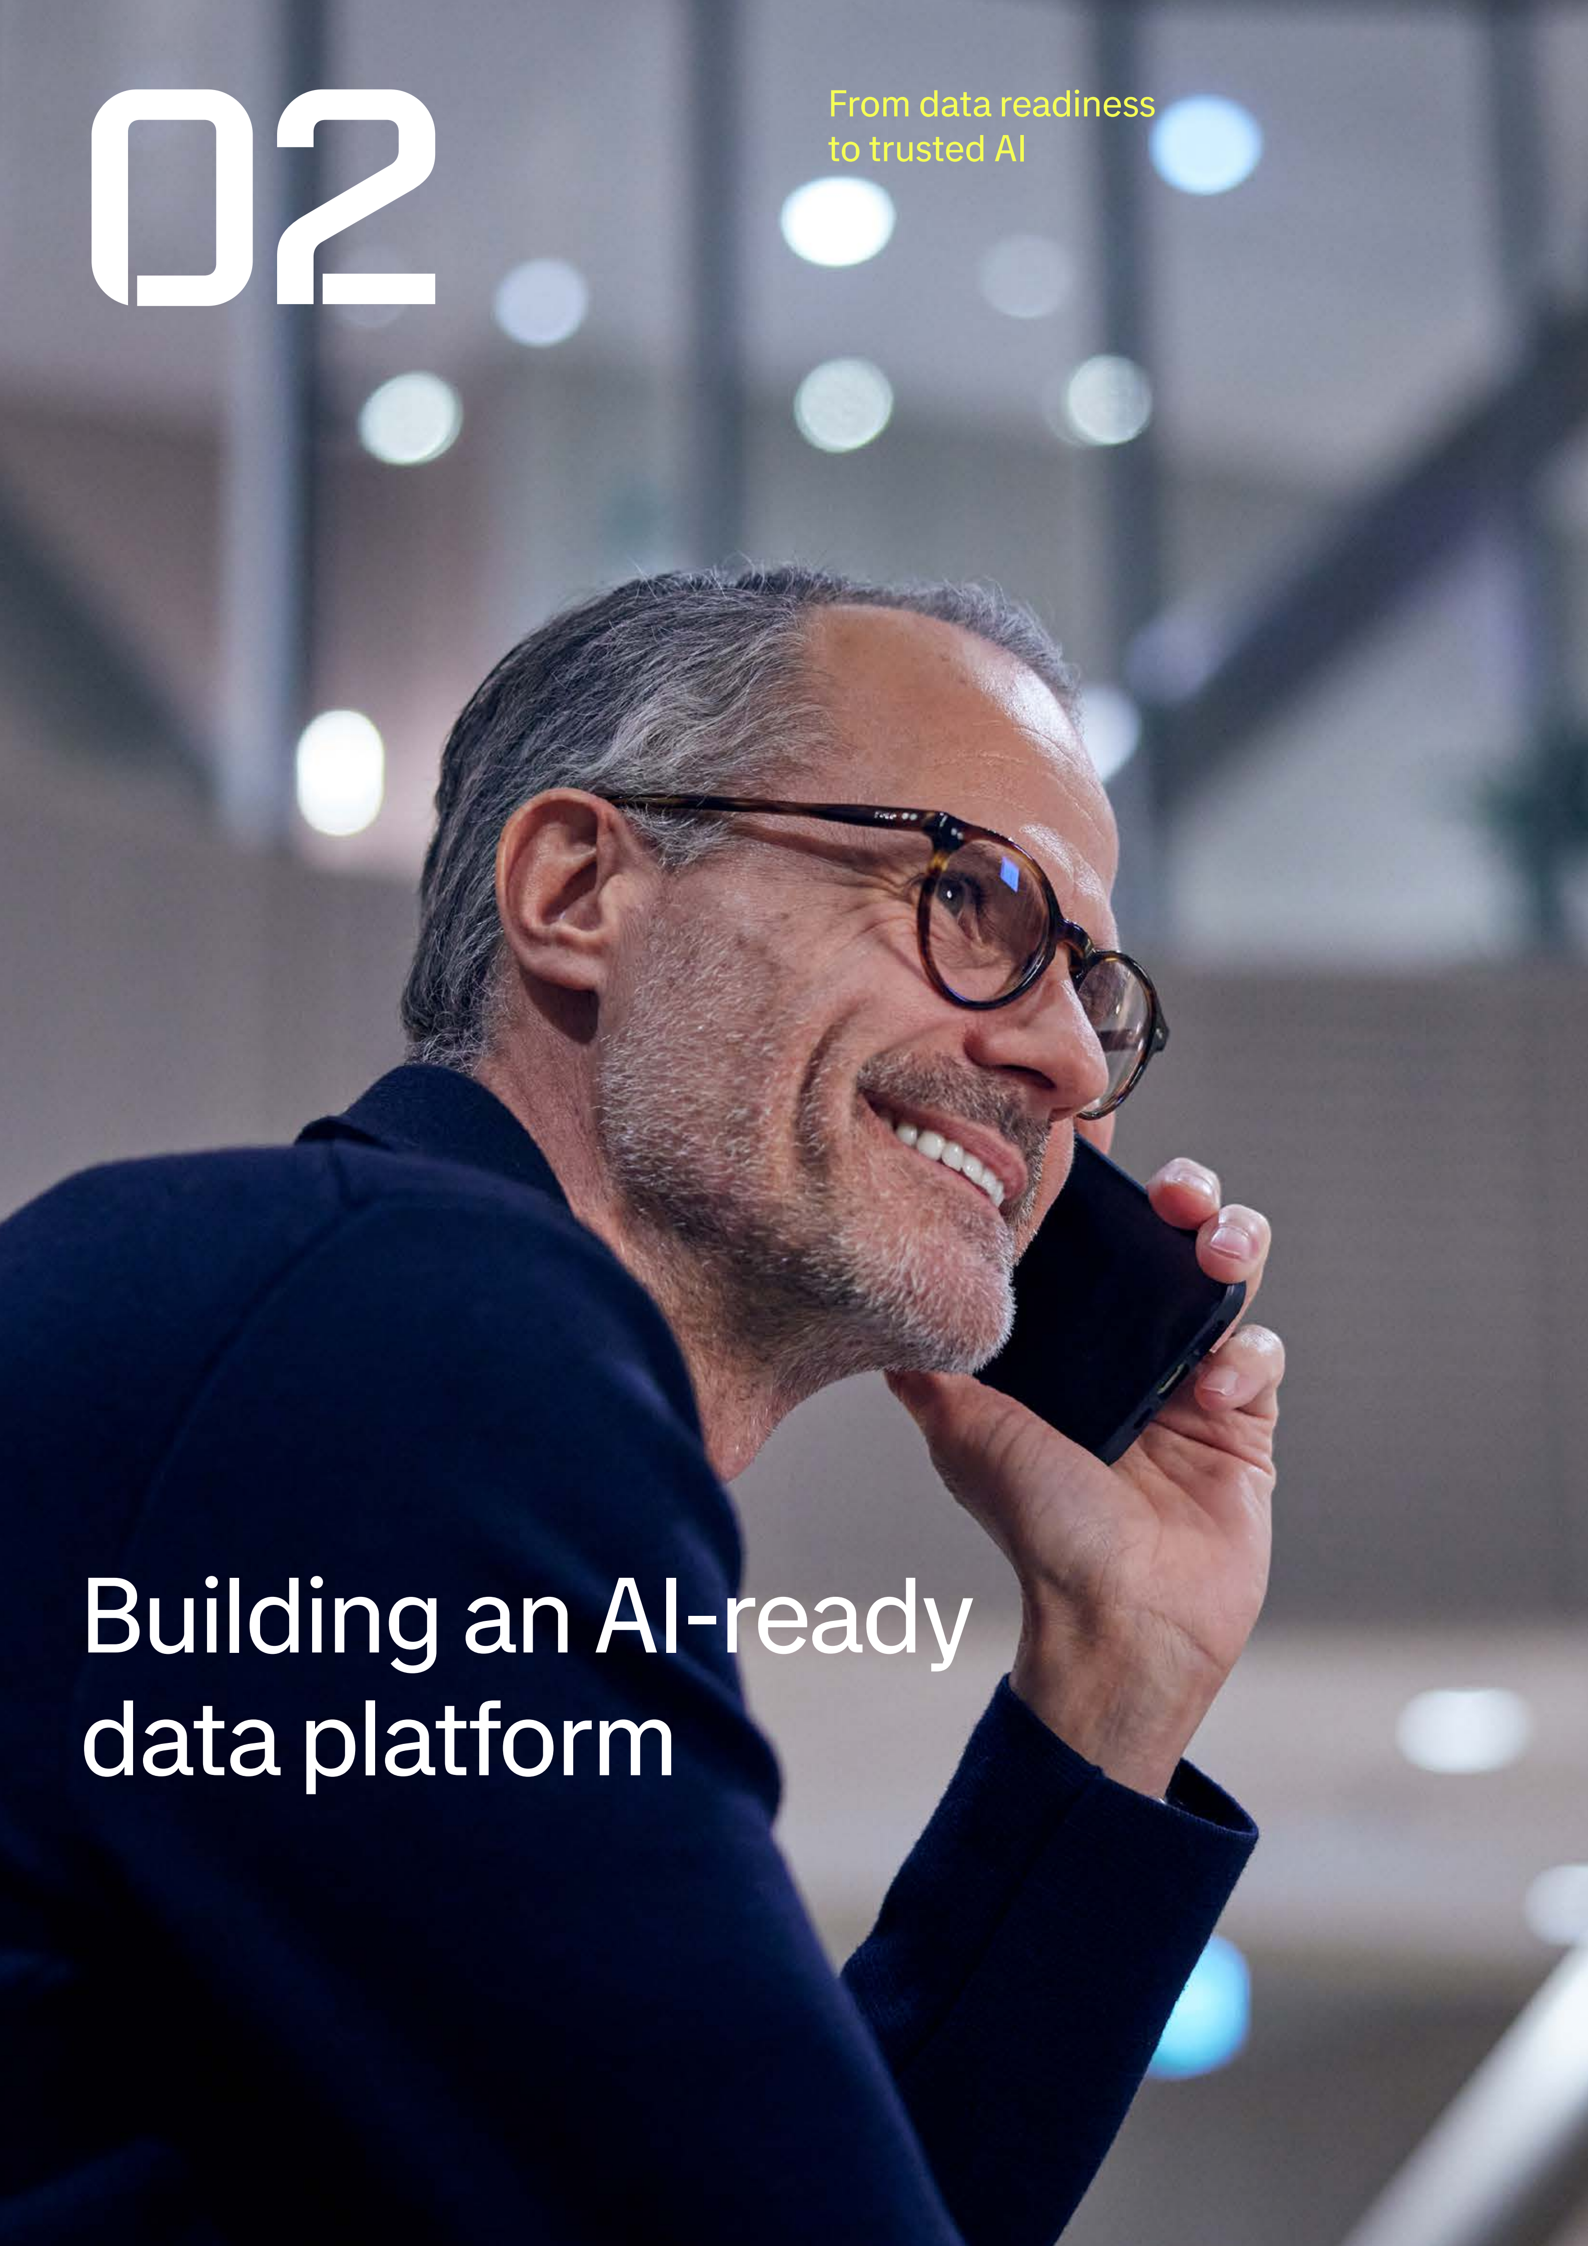

Building an AI-ready data platform

Problem

Most organizations already have significant amounts of data, but it is often fragmented across legacy systems, business applications, and departmental silos. This fragmentation makes it difficult to access, integrate, govern, and trust the information that AI solutions depend on.

At the same time, the shift from structured to massive volumes of unstructured and semi structured data (images, sensor feeds, etc.) has outpaced the capabilities of traditional data architectures. AI cannot run on infrastructure designed for dashboards. It requires a data platform that can handle speed, scale, governance, traceability, and continuous evolution. Without such a foundation, organizations struggle to create the consistent and contextual data flows needed for reliable AI.

Recommended approach

To support AI at scale, organizations need a modern data platform that can ingest, unify, govern, and process both structured and unstructured data in a reliable and flexible way. By creating a unified data environment, Tieto Tech Consulting enables consistent data flow and integration, ensuring that all relevant information is accessible and actionable. A modern data platform not only consolidates data from disparate sources but also maintains its integrity and quality, access control and transparency, providing a solid foundation for advanced analytics and AI applications.

A robust AI-ready platform should therefore provide:

- Elastic compute and storage to scale with new AI use cases
- Automated data pipelines to ensure repeatable and reliable transformation processes
- Real-time data ingestion when operational insights matter
- A modular architecture that supports innovation and reduces dependency on a single vendor model
- End-to-end governance to maintain trust and compliance
- Metadata, lineage, and provenance to trace how AI decisions are made
- Iterative growth aligned with business priorities so that each new use case strengthens the broader data foundation

A modern data platform is not just a technology upgrade. It is a strategic capability that enables organizations to deliver AI use cases faster, reduce risk, and create a stronger return on investment over time.

Case study

Cloud-native data foundation for AI

A prime example of this approach is our collaboration with a large US steel manufacturer. The company was struggling with on-premises legacy systems lacking scalability and robustness, which made it difficult to ingest disparate data sources and formats. There was no reliable long-term data storage for regulatory compliance, and the platform did not support AI initiatives. Our solution involved designing a medallion architecture – a data management approach that organizes data into distinct layers to systematically enhance data structure and quality as it moves through each layer, balancing raw data with clean, analytics-ready datasets.

The medallion architecture enabled data integration and transformation to:

- Meet varied consumption needs
- Implement enterprise-scale orchestration with Azure Data Factory for seamless data flow
- Enable incremental ingestion from multiple disparate systems
- Build a data lake on Azure Databricks
- Deliver interactive reports and dashboards using Microsoft Power BI.

The centralized data warehouse and data lake yielded several advantages. It enabled standard reporting across the enterprise and reduced infrastructure, maintenance, and scalability costs by moving to the cloud. Besides, the platform unified data from raw to curated layers and accelerated data-driven decision-making on materials and processes supported by AI models.

The platform also enabled the development of a chatbot for staff to quickly query specific data on the steelmaking process, based on the client's knowledge hub that contains extensive industry-grade research articles and metallurgy guide data in various formats.

While the process of reading and interpreting the content was previously too time-consuming for staff researchers, our solution helped:

- Enhance context understanding for the AI model by adding a retrieval step that fetches up-to-date information (Retrieval-Augmented Generation, or RAG) for more comprehensive and meaningful interactions
- Implement automated document processing pipelines to handle complex extractions from various sources
- Summarize results on materials, formulations, and processes for steel products
- Ensure enterprise security by implementing best practices for data, network, application, identity, and access security for the AI application in Azure

Executive takeaway

Legacy data environments are not sufficient to support long-term AI ambitions. A modern AI-ready data platform is essential, but its true value lies in enabling trusted, scalable, and continuously evolving data use across the business. To ensure this, business and IT stakeholders must address questions such as:

Do we have the right skills (data engineering, data science, domain expertise)?

Who will sponsor the technical implementation of the use case and remove organizational blockers?

What is the total cost of ownership (data work, infrastructure, people, operations)?

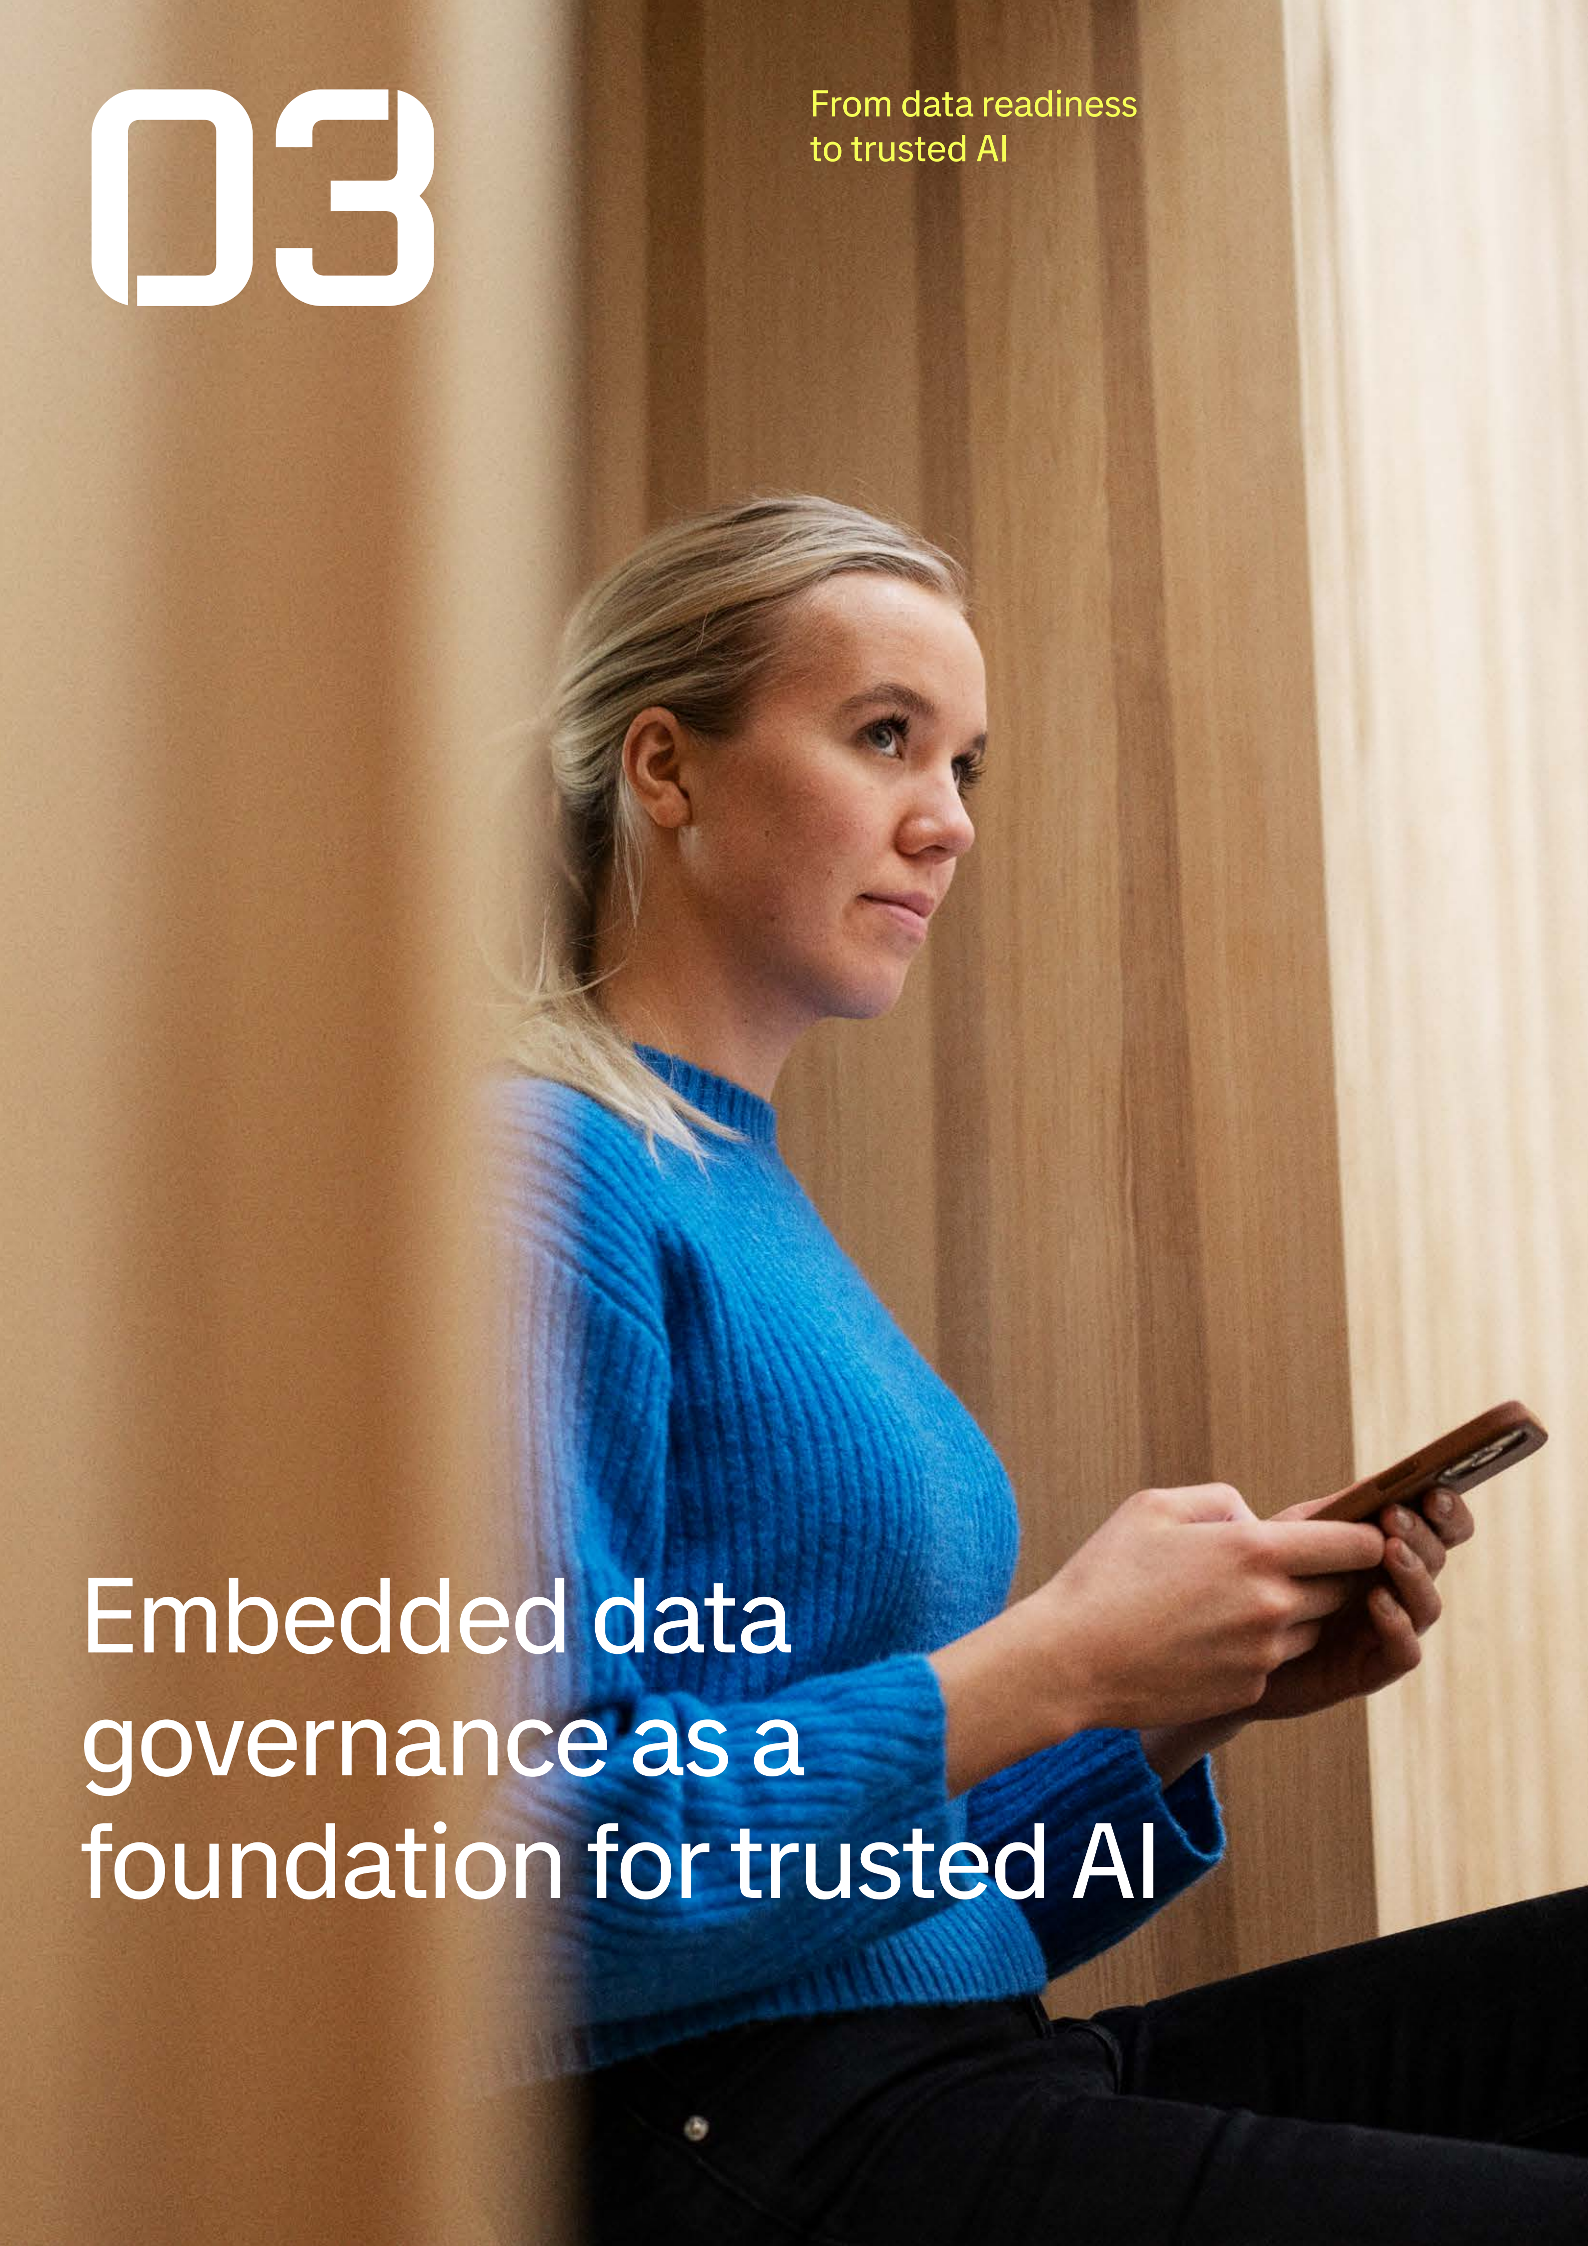

03

Embedded data governance as a foundation for trusted AI

Problem

Many AI initiatives slow down after data has been collected and integrated into a platform. The underlying issue is insufficient data quality due to fragmented master data systems, inconsistent product information, and a lack of standardization. Before data can power advanced analytics or AI, it must be strictly governed and undergo rigorous assessment, cleaning, and transformation to ensure it is accurate, timely, and has a documented lineage. Weak or fragmented governance represents a compliance and reputational risk when AI systems increasingly influence customer interactions, operational decisions, and regulatory outcomes. Without addressing these challenges, organizations risk slow development cycles, weak model performance, and reduced trust in AI systems.

Recommended approach

A structured and continuous approach to data quality is required, grounded in strong governance and data lifecycle management. **Data governance** defines clear roles, ownership, and usage rights. **Data management** ensures data is properly handled throughout its lifecycle. Both are essential to ensure trust in data, address ethical considerations, and comply with regulatory frameworks such as the EU AI Act or the General Data Protection Regulation (GDPR).

A robust approach to data governance therefore requires clear executive ownership and must be anchored in the organization's overall risk and compliance framework. Responsibility for data governance should be explicitly assigned at C-level, with defined mandates, escalation paths, and decision rights that reflect the business impact of AI-driven decisions.

Key governance elements include:

- **Ownership and stewardship:** Define accountable data owners and stewards and enforce policies across the entire data lifecycle, including model access and development.
- **Data lineage and provenance:** Track the origin, transformations, and context of data.
- **Access management:** Implement clear rules to ensure employees only access the data they are authorized to use.
- **Standardization:** Implement consistent formats, units, definitions, and taxonomies.
- **Lifecycle management:** Manage updates, versioning, and validity.

Our approach to data governance starts with a structured data quality assessment where we

- Evaluate the current state and completeness
- Identify inconsistencies, duplicates, outdated formats or missing context
- Define corrective actions

Based on this assessment, iterative actions to improve data quality must be defined. It is important to understand that data governance and data management are not a one-time effort but an ongoing process. Data is controlled, traceable, and auditable for one use case. Once that system proves valuable, data governance can be expanded to the next area. For this purpose, strong data governance and data management practices must be established and aligned with business strategy, to break down data silos and build a well-managed data foundation step-by-step.

Executive takeaway

Data quality is not a one-time exercise. It is an ongoing capability that requires governance, accountability, and continuous improvement. Organizations that embed these practices create the trust, transparency, and control needed to scale AI responsibly. Here it is important to answer:

How are clear accountability and governance structures established within the organization?

How are data quality standards defined, agreed, and enforced across business and IT?

Are incentives and KPIs aligned so that teams are rewarded for maintaining data quality, not just delivering features?

No data, no differentiation. The real AI value comes from your data.

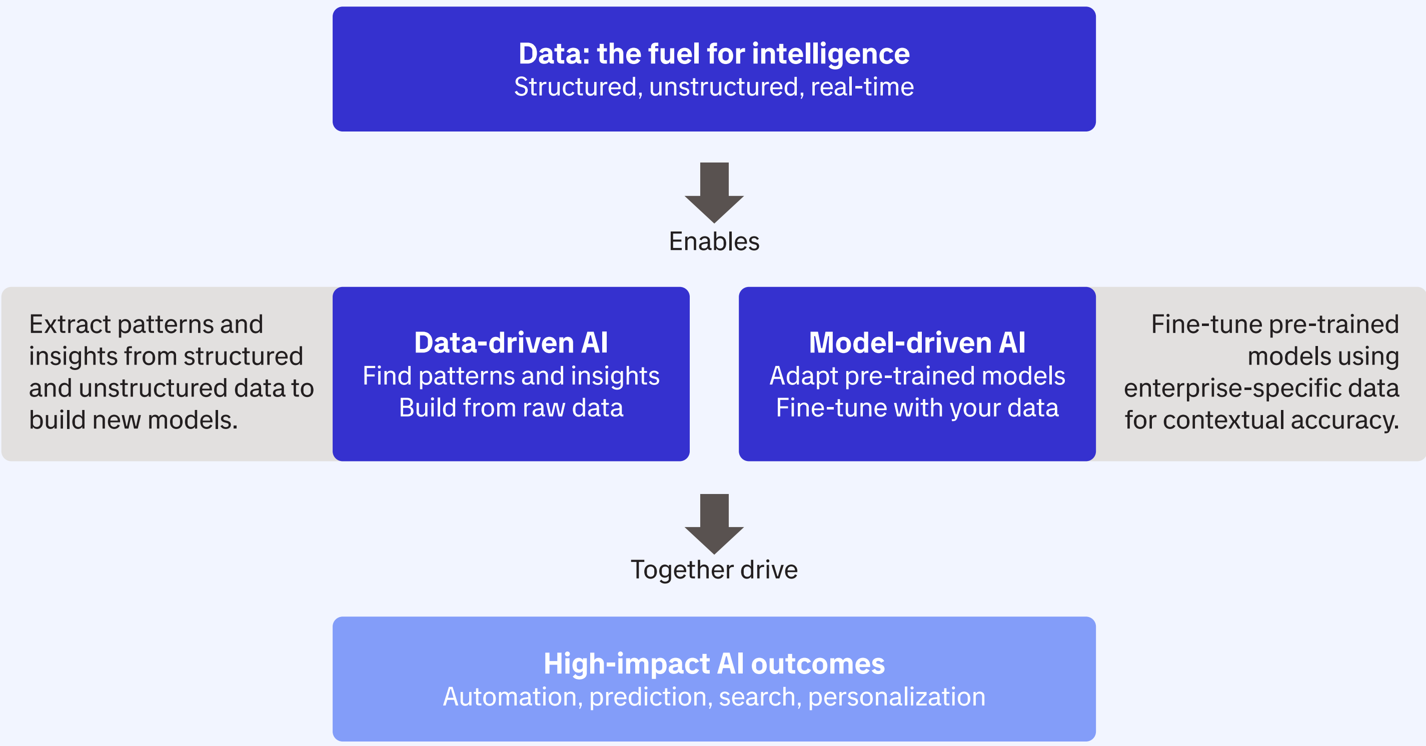

Data driven AI relies primarily on large volumes of high quality data to learn patterns, making data availability, accuracy, and context critical for performance. In contrast, model driven AI begins with predefined rules, domain knowledge, and expert designed structures, allowing models to operate effectively even with limited data. Together, these approaches demonstrate how organizations can balance empirical learning with structured reasoning to build robust, reliable, and explainable AI systems. Common to both approaches is a solid data foundation, grounded in data governance and data management.

04

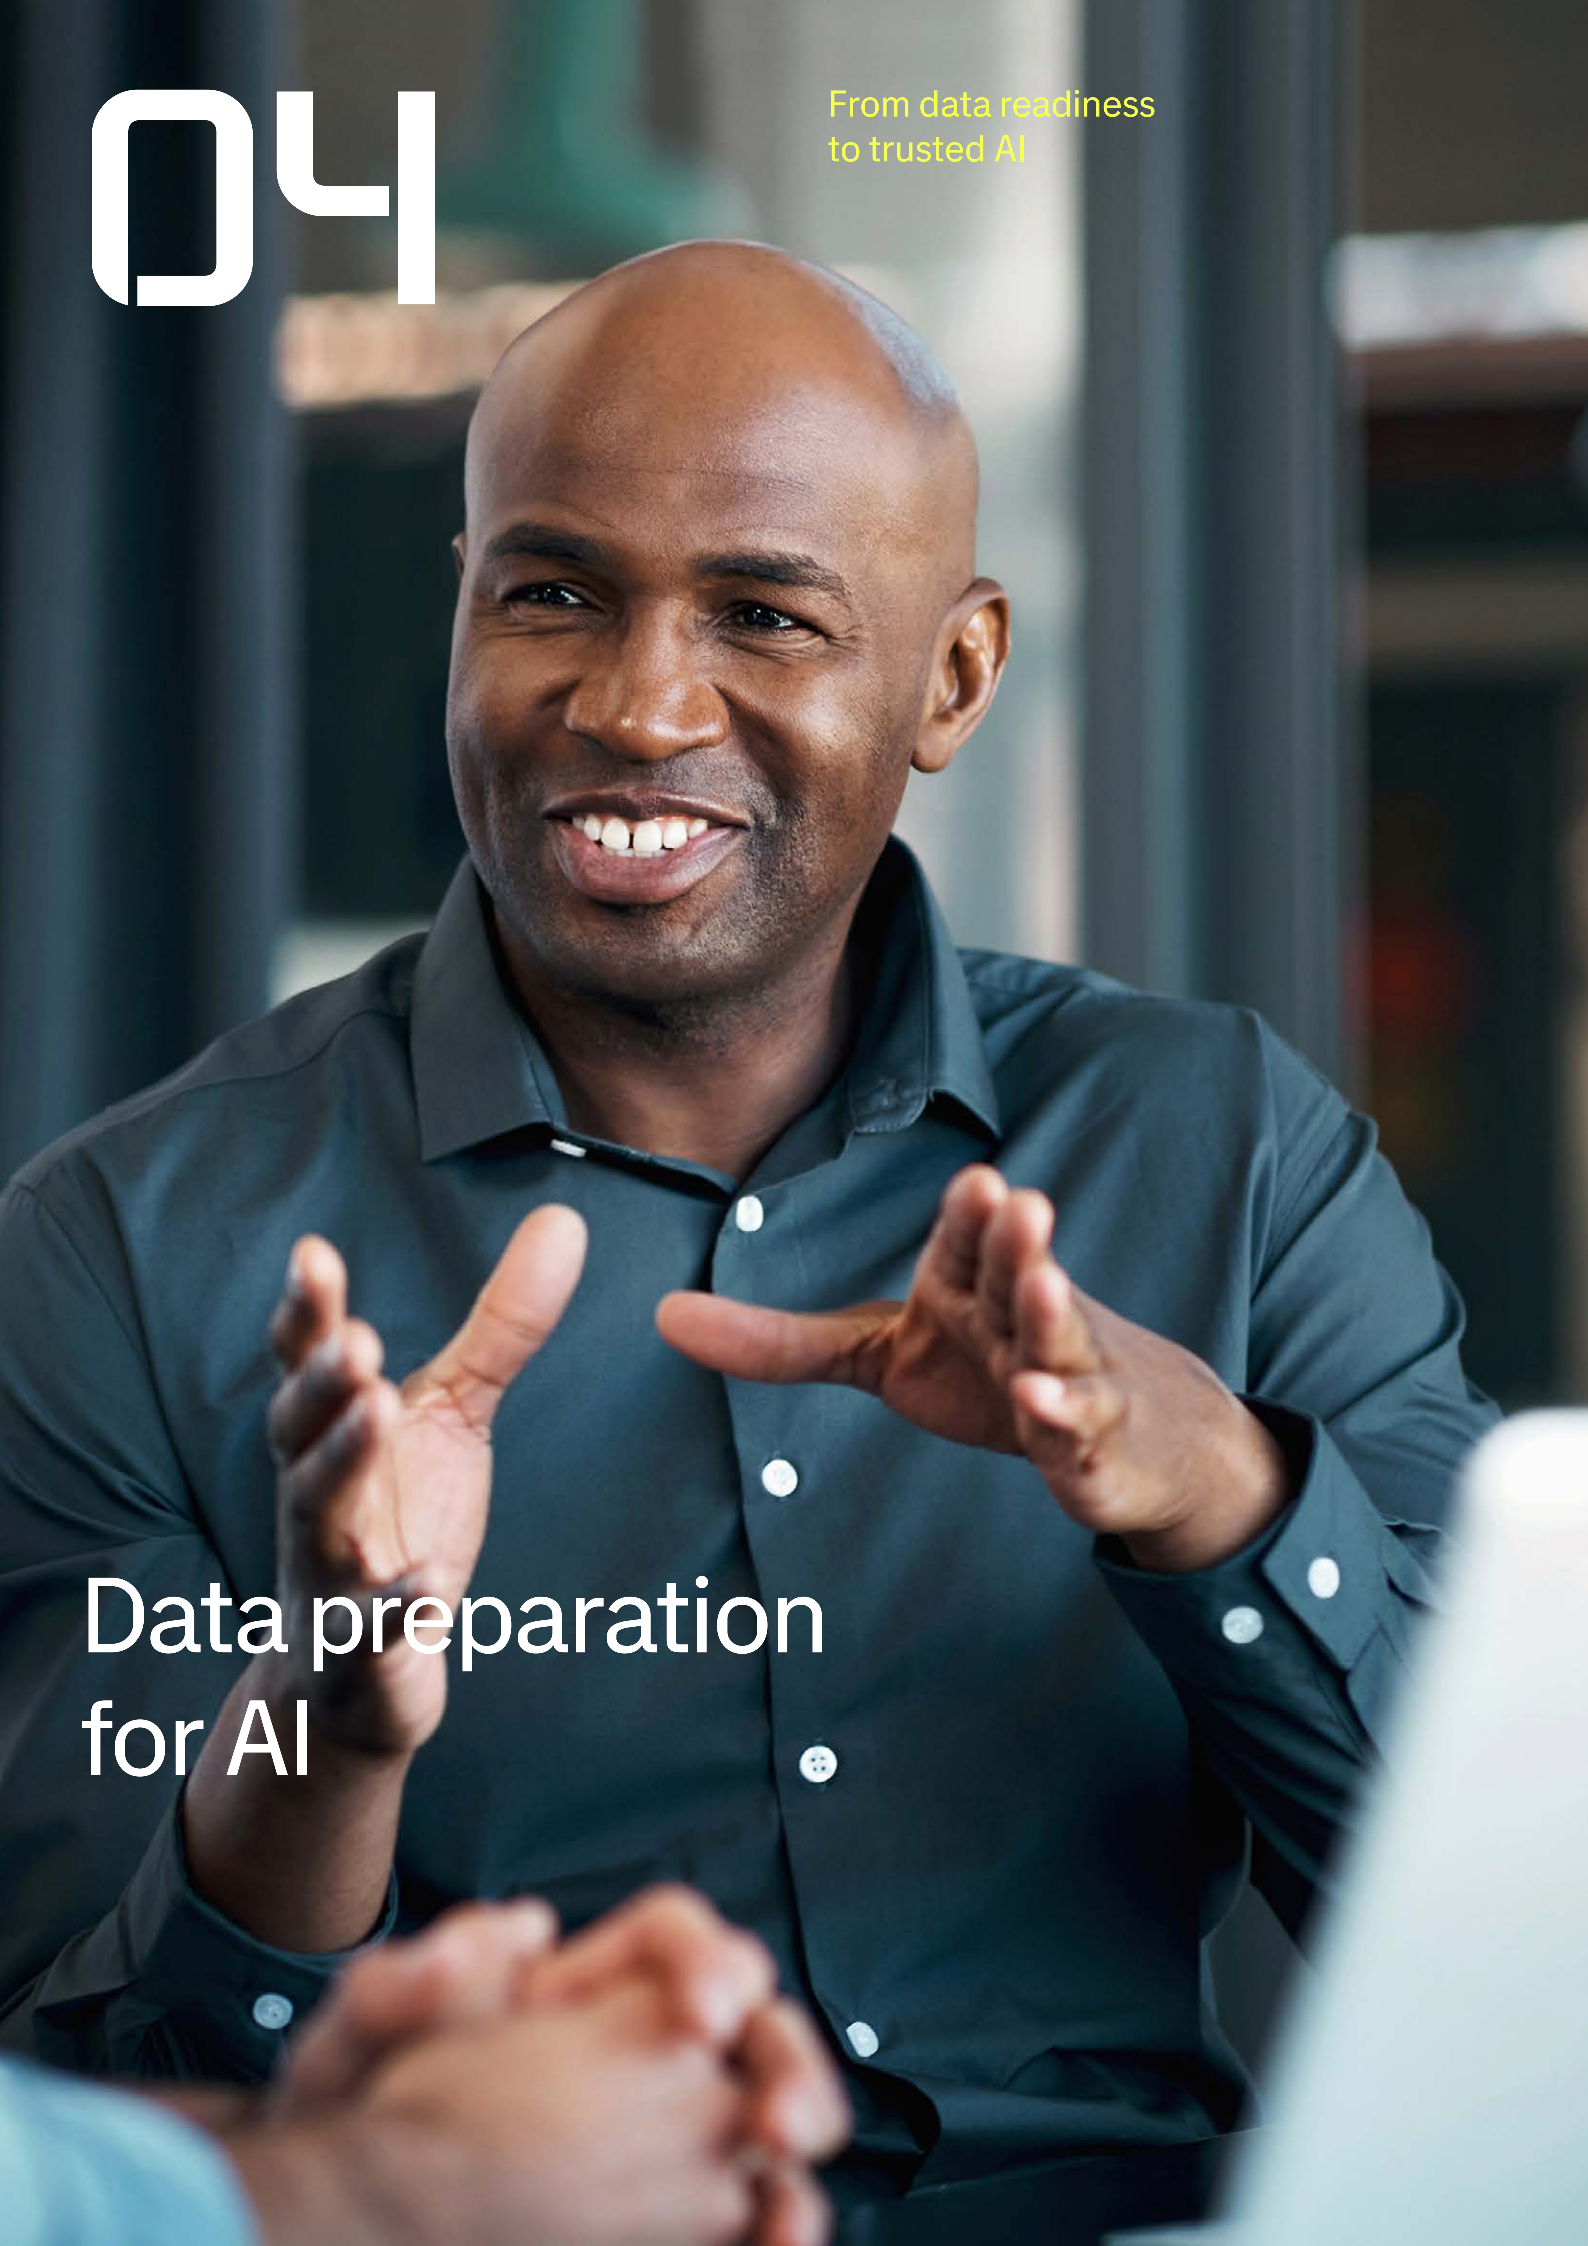

Data preparation for AI

Problem

Without data quality assessment and the establishment of data governance, downstream activities for AI become inefficient and AI performance unreliable. Without this foundation, data preparation becomes ineffective and risks embedding errors directly into the training data.

AI workloads rely on clear business metadata and semantic model structures that make the data understandable and reusable. This is particularly important when human annotators or automated systems need sufficient business context to assign the right labels and transform and interpret the data correctly.

The challenge becomes even more visible with unstructured data such as images, video, and text. Without enriched metadata, datasets remain opaque and organizations struggle to reuse or operationalize data assets at scale. However, cleaning, transforming, and labelling the data can become a bottleneck for any AI initiative, especially if there is a lot of data to annotate and few resources available for this, or if this cannot be done automatically.

From a business perspective, neglecting these prerequisites results in:

- Delayed time-to-value due to rework
- Increased operational costs for cleaning and re-labelling
- Model outputs that the business cannot trust
- Lower user adoption due to accuracy issue.

Further, domain and process experts who truly understand the data and processes that generate it are crucial for AI model development. Not understanding every detail of the data and process might lead to AI models which perform well in testing but fail miserably in production because data gets updated in ways that cannot be traced after the fact.

Recommended approach

Metadata is a foundational element of AI readiness. Organizations should establish a structured approach to metadata and labelling so that data can be interpreted, governed, and reused consistently.

Here are some of the recommended practices:

- Ensure early engagement of subject matter experts to understand the data and the processes generating it.
- Implement a robust metadata strategy, ensuring datasets include schema information, business terms, semantic tags, and lineage.
- Use data catalogs and business glossaries to make metadata searchable and consistent across domains.
- Wherever possible, automate metadata harvesting.
- Apply phased or semi-automated labelling approaches to reduce manual work.
- Preserve business meaning in labels by linking annotation categories to metadata tags.

High-quality metadata and accurate labelling enhance:

- **Traceability** through lineage and provenance metadata that document how datasets were created, transformed, and used.
- **Security and access controls,** where metadata classification enables role-based controls and compliant usage.
- **Model transparency,** as metadata gives visibility into assumptions, input sources, and semantic definitions used during model training.

Well-labelled datasets also enable continuous validation and monitoring, ensuring that AI models remain aligned with business requirements and do not drift away from intended behaviour.

Case study

A phased approach for supervised learning

We implemented a two-step approach in a project for a manufacturing client that was facing challenges with image labelling due to a lack of resources. The project sought to provide automatic detection of product quality control: “good” products can be processed further, while root cause of “bad” ones must be analyzed to improve production.

In the first step, we focused on anomaly detection to solve the main business problem and have a high-level sorting of the products between “good” and “bad” ones. The system was trained to recognize correct images, while all types of “bad” ones were labelled as “anomaly”, regardless of the root cause. When an anomaly was detected, the system would store the corresponding images in a dedicated folder, thus reducing the amount of data and simplifying the labelling process. This already enabled “good” pallets to be processed, while “bad” ones were recognized early enough. As a second step, only “bad” images were labelled, which then served as a basis to train a classification model more quickly and efficiently. The detailed classification of the detected anomalies ultimately enabled prescriptive responses.

This step-by-step approach allowed us to quickly detect anomalies and then move on to detailed classification, always keeping the business problem in focus. The phased approach proved essential for successful implementation and user adoption of the model.

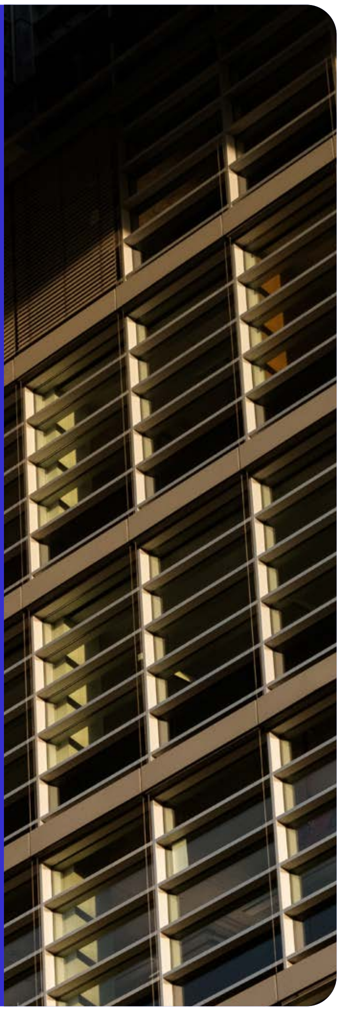

Executive takeaway

AI depends on data that is not only available but also well described, semantically clear, and labelled with business relevance. Strong metadata and thoughtful labelling practices are essential for creating reliable, scalable AI solutions and can be ensured by addressing the following topics:

Who is responsible for defining, maintaining, and approving metadata and labels as the business evolves?

How do we ensure that metadata and labels stay synchronized with reality when processes, products, or regulations change?

Do tools and processes make metadata and labels discoverable and mandatory, not optional documentation?

05

Model training and deployment: balancing automation with human oversight

Problem

In many organizations, model training is treated as a one-off technical exercise, where data scientists build a model, deploy it, and move on. In reality, models degrade quickly when exposed to changing business conditions, evolving data patterns, and real-world operational complexity. This challenge is particularly significant in dynamic environments such as industrial automation, computer vision, and multi-site operations, where data drift and changing conditions can quickly reduce model performance.

Furthermore, AI workloads introduce additional requirements around governance, compliance, and auditability. Business leaders and end users are not willing to rely on "black box" models that degrade unnoticed. They need scalable and auditable models that are aligned with operational KPIs.

Without a structured operational framework, organizations face:

- Model drift leading to inaccurate predictions
- High maintenance costs due to manual retraining
- Slow time to value and inconsistent performance across business units
- Reduced trust from regulatory, operational, and business stakeholders

Recommended approach

A successful machine learning operations (MLOps) approach transforms machine learning from isolated technical experimentation into a scalable, business driven capability. Rather than treating models as one off deliverables, organizations should operationalize ML as a repeatable, end to end lifecycle that includes development, deployment, monitoring, retraining, and governance.

Our recommended approach includes:

- Standardized MLOps framework: establish a unified operating model that seamlessly integrates data engineering, model development, deployment, and monitoring on a single platform to reduce complexity and time to value.
- Automation and reliability at scale:
    - Version control for code, data, and models to ensure transparency and auditability
    - Automated training and deployment pipelines to accelerate releases and reduce manual risk
    - Clear stage transitions (development → testing → production) to support enterprise grade governance
- Built-in governance and performance control to ensure reproducibility and regulatory compliance, maintain model quality over time, and proactively manage performance degradation and business risk
- Incremental adoption aligned to business priorities: start with automation and traceability for high value use cases, then evolve toward full lifecycle management as maturity increases.

This pragmatic approach helps organizations scale machine learning reliably while improving efficiency, reducing risk, and maximizing business impact.

MLOps cycle

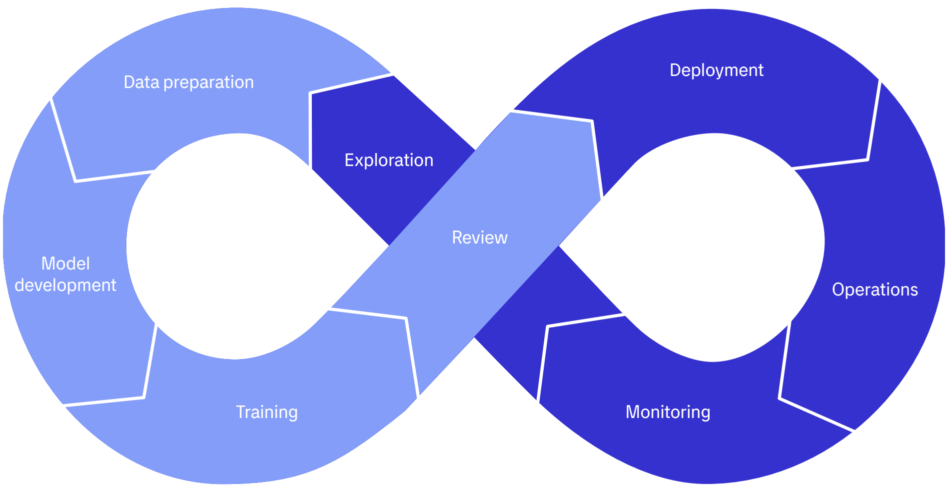

Case study

Sustaining AI performance through MLOps

In a recent computer vision initiative for a leading automotive sector client, we addressed the challenge of maintaining state-of-the-art models in a dynamic operational environment by implementing a comprehensive MLOps framework.

The objective was to deliver reliable, real-time visual analytics across multiple sites, where changing conditions and new data streams required models to be continuously updated. Traditional approaches, where models are trained once and deployed statically, proved insufficient, as model performance would degrade over time due to data drift and evolving operational requirements.

We solved the problem by implementing an end-to-end MLOps pipeline that automated the entire lifecycle of our computer vision models. This included:

- **Automated data ingestion and preprocessing:** New image and sensor data from operational systems are continuously ingested and preprocessed, ensuring that the training data always reflect the latest realities.
- **Continuous training and experiment tracking:** Using a standardized pipeline, models are retrained on fresh data at regular intervals. All experiments, hyperparameters, and results are tracked to ensure reproducibility and facilitate rapid iteration.
- **Model validation and governance:** Each new model version undergoes rigorous validation against hold-out datasets and business KPIs. Only models that meet predefined performance thresholds are promoted to production, ensuring reliability and transparency.
- **Seamless deployment and rollback:** Deployment is fully automated, with new models rolled out to production environments with zero downtime. In case of performance degradation, the system can automatically revert to a previous stable version, minimizing operational risk.
- **Real-time monitoring and alerting:** Model performance is continuously monitored in production. Automated alerts notify the team of any anomalies, triggering retraining or investigation as needed.
- **Human-in-the-loop oversight:** While the retraining and deployment processes are highly automated, the client retains full control at every stage. At any point, authorized users can intervene by pausing automation, reviewing model performance, approving or rejecting retraining cycles, or manually triggering model updates. This ensures that expert judgement and business context are always incorporated, particularly in critical or exceptional scenarios.

Executive takeaway

The challenge is not simply to train a model but to sustain its business value over time. MLOps enables organizations to scale AI with greater reliability, agility, transparency, compliance, and human oversight. Following questions help address long-term business value for MLOps:

How is human oversight maintained for high-impact or high-risk decisions?

Does our MLOps setup allow us to scale AI reliably across use cases without reinventing controls each time?

How do we continuously measure business value, not only technical metrics (e.g. accuracy vs. impact on KPIs)?

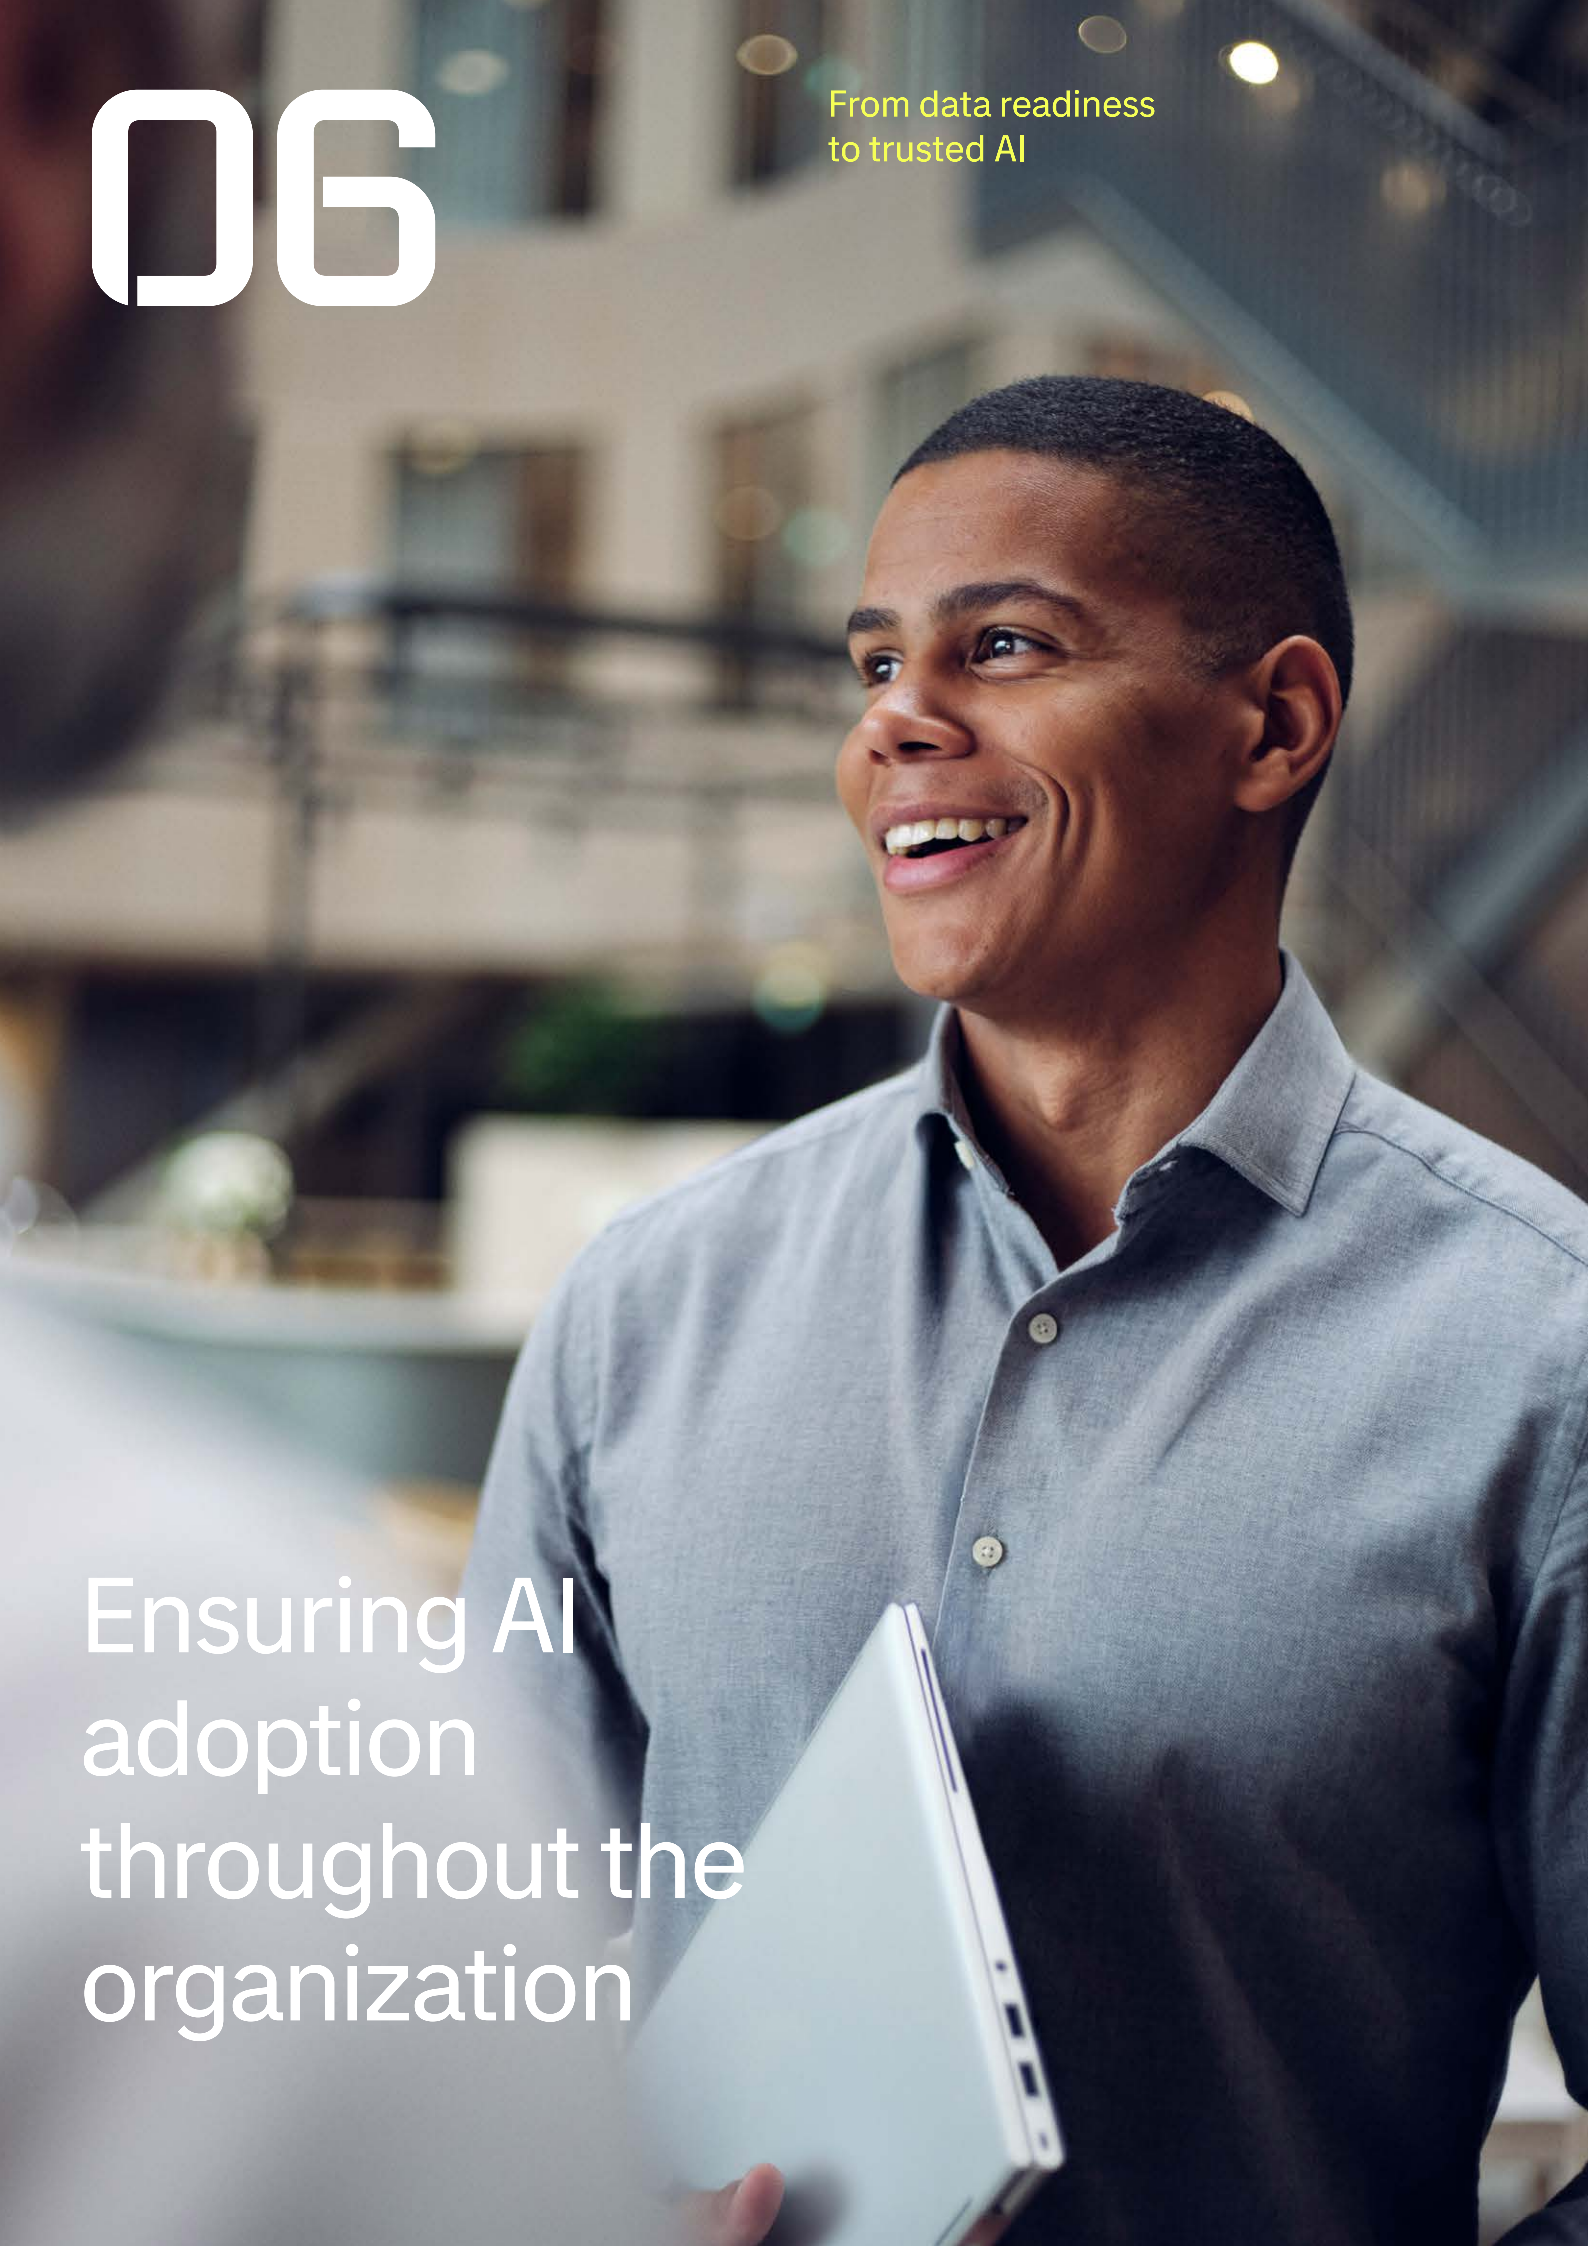

06

Ensuring AI adoption throughout the organization

Problem

Successful AI adoption depends as much on people and processes as it does on technology. Even technically strong solutions fail to create business impact if users do not trust them, do not understand them, or do not see how they fit into daily work.

Common blockers include:

- Low data and AI literacy across the workforce, creating resistance to using the new technology
- Lack of trust due to limited transparency or insufficient understanding of the support provided by AI
- Inadequate change management, causing promising proofs of concept to fail at scale

Without focused adoption measures, organizations risk building technically capable solutions that never achieve meaningful business usage.

Recommended approach

Effective adoption starts long before rollout. Organizations need structured change programs, using models like ADKAR (Awareness, Desire, Knowledge, Ability, Reinforcement) that focus on transparency, communication, and incentives:

- **Awareness** that AI changes how decisions are made, not just how tasks are executed.
- **Desire** through early wins and visible executive sponsorship. This can be fostered by starting in an area with high data quality due to standardization, such as invoices or customer orders.
- **Knowledge** via targeted upskilling on data literacy and AI outputs but also via early objection handling and address of user concerns, especially around job impact, data usage, and accountability. This is strongly supported by clear communication of the benefits and limitations of AI solutions.
- **Ability** by embedding AI into actual workflows rather than running it as a parallel tool and by enabling continuous knowledge transfer, allowing teams to expand AI usage autonomously. Here it is also important to consider “human-in-the-loop” for cases that fall below a defined quality threshold to enhance user trust and confidence.
- **Reinforcement** through KPIs that measure adoption and business impact, not just deployment.

Trust in AI can be established when teams start to understand data quality, model outputs, and their implications for decisions. This is ensured by cross-functional collaboration between business, IT, and data teams.

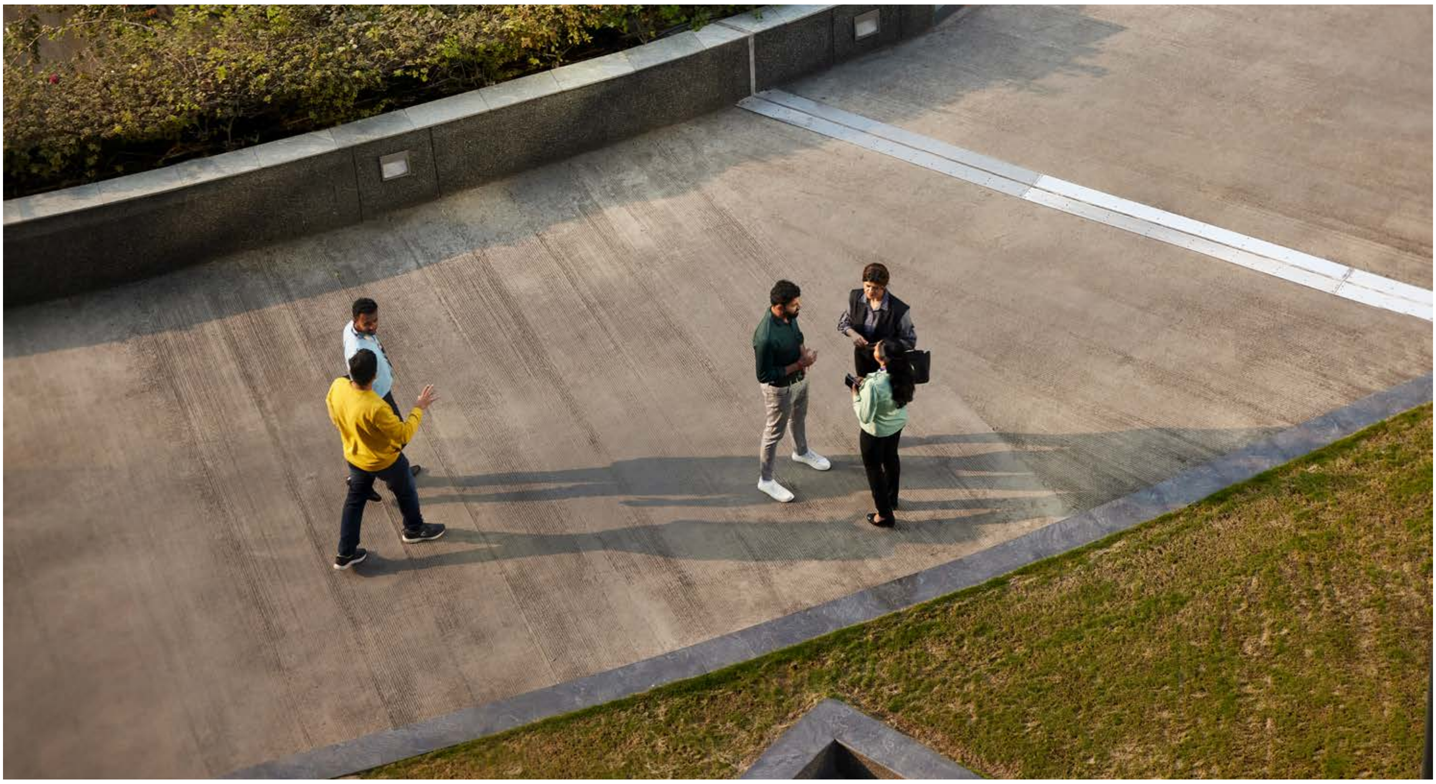

Case study

Securing AI adoption at scale

The scope of our customer case included 15 countries with up to 110 sites. The goal was to automate incoming orders in different formats flying in via email. Extracting the relevant data and creating the order in SAP was an entirely manual process at our customer.

Our approach encompassed not only the definition of a modern target IT architecture but also change management. We analyzed business needs to ensure that the proof of concept (PoC) provided compelling evidence of the AI technology's added value to the business, which was crucial in gaining user trust and acceptance. A joint rollout of the solution was conducted including a continuous know-how transfer from Tieto Tech Consulting as partner to the customer.

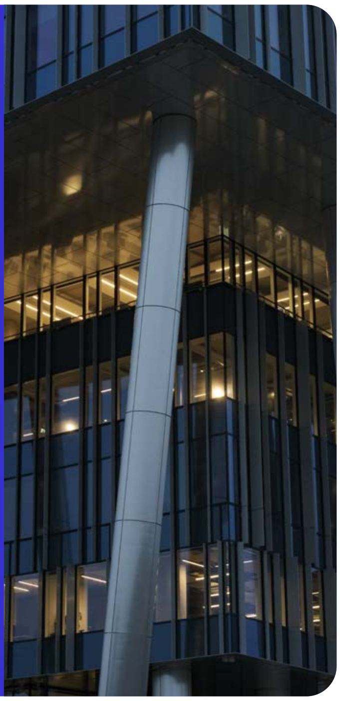

Executive takeaway

AI fails when users do not understand, trust, or see value in AI. AI adoption is an organizational shift that succeeds only when people are ready, confident, and willing to use AI to its full potential. Change management should address the following questions:

At what point should we stop, pivot, or scale the initiative?

Are business users trained to understand data limitations and to challenge AI outputs when needed?

How do we address increased scrutiny from regulators, customers, or employees?

Conclusion and outlook: data readiness in the age of agentic AI

AI success does not begin with the model. It begins with the readiness of the data, the strength of the governance framework, and the organization's ability to operationalize AI at scale. Companies that treat data readiness as a strategic capability, and not only as a technical afterthought, are better positioned to move faster, reduce risk, and generate measurable business value from AI.

The six steps outlined in this paper show that AI readiness is not achieved through a single platform decision or isolated pilot. It requires a deliberate progression: starting with business value, establishing a scalable data foundation, governing data with clarity and accountability, preparing and enriching datasets for AI use, operationalizing models through MLOps, and enabling adoption across the organization. When these elements come together, AI becomes more than experimentation. It becomes a repeatable engine for performance, resilience, and innovation.

This is even more important for the implementation of **agentic AI systems** designed to autonomously plan, coordinate, and execute multi step workflows, often across multiple data sources, models, and systems. In traditional AI use cases, data issues may surface as degraded model performance or delayed insights. In agentic systems, the same issues can lead to uncoordinated actions, inconsistent outcomes, or loss of control, especially when multiple agents interact across processes and domains, in real time and with limited human intervention. Agentic AI significantly raises the bar for data readiness and turns the recommended approach into prerequisites:

- Governance can no longer be an overlay after deployment; it must occur in real time and be embedded into data architectures by design.
- Data quality can no longer be assessed periodically; it must be continuously monitored and enforced.
- Metadata can no longer be treated as disconnected attributes. A semantic layer must provide the links between business meaning and machine-readable form for AI agents to act upon.
- Transparency is no longer optional, as organizations must be able to explain not only model outputs but agent decisions and actions across entire workflows.

Organizations that will lead with agentic AI are those that first build trust in their data. By investing in data readiness now, companies create the conditions for faster deployment, more reliable outcomes, stronger compliance, and sustainable user adoption. Organizations that treat data as a well governed, shared enterprise asset will be able to deploy AI agents with confidence, control, and measurable impact. In that sense, data readiness is not only the gateway to AI success but also the essential foundation for long-term competitive advantage.

About Tieto Tech Consulting

Tieto Tech Consulting provides design-led, data-centric, and AI-powered digital engineering & consulting services to enterprises worldwide. With a strong heritage of Nordic quality, local presence, and global-scale operations, we are backed by 8,500+ IT professionals. Our team elevates people-first experiences and builds tailored digital solutions, cloud, BI, and AI systems.

We integrate modern data platforms and unify diverse enterprise software solutions into scalable digital ecosystems, helping you efficiently scale and accelerate your business while making it smarter with purposeful technologies. Our clients trust us for advanced data and analytics skills, top industry-specific expertise, specialized software R&D, and the capabilities of a multinational team. We are part of Tieto, a leading technology company with an annual revenue of approximately EUR 2bn.

Contact us

Anne Marchal
Senior Solution Consultant
Tieto Tech Consulting

anne.marchal@tieto.com