INOVIQ
TECHNOLOGIES
Engineering Intelligent Industry

# ARTIFICIAL INTELLIGENCE IN MANUFACTURING AND PROCESS INDUSTRIES

## A PRACTICAL CUSTOMER GUIDE

Proven applications, quick wins, standards, and data sovereignty and protection
Companion to the INOVIQ Digital Transformation Primer and Maturity Self-Assessment

**Contents**

1 Executive summary
2 AI terms and types used in industry
3 Where AI delivers proven value
4 Choosing the right use cases
5 Reference architecture for industrial AI
6 Ten quick wins
7 Generative AI in industrial settings
8 Data readiness
9 Data sovereignty and data protection
10 Governance, safety and standards
11 The AI life cycle and MLOps
12 Roadmap, KPIs and next steps

# 1 EXECUTIVE SUMMARY

Artificial intelligence (AI) is already proven in manufacturing and process plants. Computer vision inspects products faster and more consistently than manual checks; machine learning (ML) detects equipment problems weeks before failure; models estimate quality in real time between lab samples; and generative AI helps people find answers in thousands of pages of procedures and manuals.

Success rarely depends on the most advanced algorithm. It depends on a clear business problem, good data in context, people who trust and use the result, and careful handling of the data itself: where it is stored, who can see it, and whether it is used to train someone else's model. This guide covers all four, with ten practical starting points.

**Key messages**

- **Start with proven use cases** that address daily losses: inspection, equipment health, quality, energy, knowledge.
- **Advise first, automate later:** begin with recommendations a person approves; automate only within tested, bounded limits.
- **Data decides the outcome:** contextualized, good-quality data matters more than the choice of model.
- **Keep control of your data:** choose deployment patterns and contracts that respect sovereignty, privacy and process know-how.
- **Govern AI like any critical system:** use ISO/IEC 42001, risk management and life cycle practices from the first pilot.
- **Keep AI outside safety functions:** engineered protection and safety instrumented systems stay deterministic and independent.

# 2 AI TERMS AND TYPES USED IN INDUSTRY

ISO/IEC 22989 defines AI concepts and terminology, and ISO/IEC 23053 describes how machine-learning systems are built. The terms below are those most used in industrial projects.

| Term | Meaning in an industrial context |
|---|---|
| **Artificial intelligence (AI)** | Systems that perform tasks normally requiring human judgment, such as recognizing defects, predicting outcomes or answering questions |
| **Machine learning (ML)** | AI that learns patterns from data rather than following hand-written rules |
| **Supervised learning** | Learns from labeled examples, for example images marked good or defective, or lab results paired with process data |
| **Unsupervised learning** | Finds structure without labels, such as clusters of operating states or unusual behavior (anomaly detection) |
| **Reinforcement learning** | Learns actions by trial and reward, usually in a simulator first; used in some control and scheduling problems |
| **Deep learning** | ML using multi-layer neural networks; the basis of modern computer vision and language models |
| **Computer vision** | AI that interprets images and video: defect detection, measurement, reading labels, safety zones |
| **Soft sensor (inferential sensor)** | A model that estimates a hard-to-measure variable, such as a quality property, from available process measurements |
| **Digital twin** | A digital representation of an asset or process kept in sync with real data (see ISO 23247 for manufacturing) |
| **Generative AI, large language model (LLM)** | AI that produces text, code or images; LLMs understand and write natural language |
| **Retrieval-augmented generation (RAG)** | An LLM that answers using your own approved documents, citing them, rather than general knowledge alone |
| **Edge AI (inference at the edge)** | Running trained models on devices near the machine for fast response and to keep data local |
| **MLOps** | Practices to deploy, monitor, version and retrain models reliably, like DevOps for ML |
| **Data drift, concept drift** | Changes in input data or in the relationship being modeled that reduce accuracy over time |
| **Explainability (XAI)** | Showing why a model gave a result, so people can check and trust it |
| **Human-in-the-loop** | A person reviews or approves AI output before action is taken |

### Levels of autonomy

| Level | What AI does | Who acts | Typical use |
|---|---|---|---|
| **1 Inform** | Shows insights, trends and alerts | People decide and act | Dashboards, anomaly alerts |
| **2 Advise** | Recommends an action with its reasoning | People approve and act | Setpoint advice, maintenance advisories |
| **3 Act with approval** | Prepares the action; executes after approval | People approve; system acts | Schedule changes, work order creation |
| **4 Act within limits** | Acts automatically inside tested, bounded limits | System acts; people supervise and can override | Advanced control, vision reject gates |

## 3 WHERE AI DELIVERS PROVEN VALUE

AI applications now span the whole value chain. Those listed here are in routine industrial use, not research. They share a pattern: a frequent decision or inspection, enough data to learn from, and a measurable cost of getting it wrong.

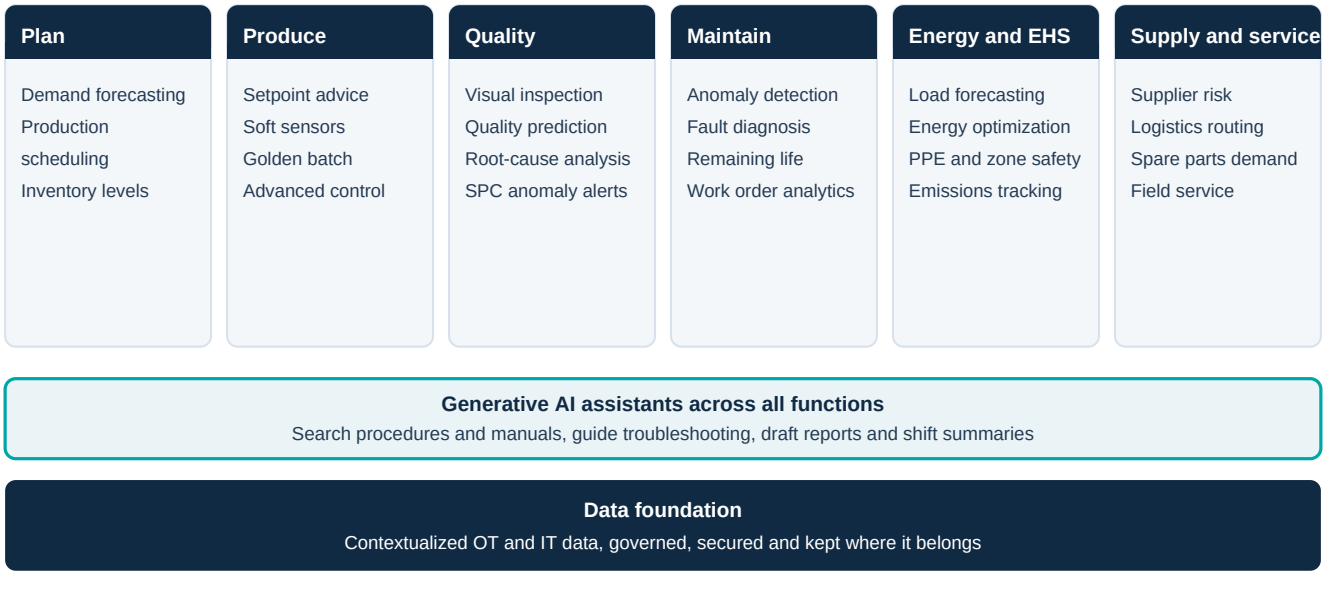


**Figure.** Proven AI applications across the manufacturing value chain, resting on a common data foundation.

| Area | Proven application | AI technique | Typical outcome |
|---|---|---|---|
| **Quality** | Automated visual inspection of surfaces, assemblies, labels, fill levels | Computer vision, deep learning | Consistent detection; fewer escapes; inspection of every unit |
| **Quality** | Real-time quality prediction between lab samples | Soft sensors, regression | Earlier correction; less off-spec product |
| **Maintenance** | Equipment anomaly detection and fault diagnosis | Unsupervised and supervised ML | Fewer unplanned stops; planned repairs |
| **Process** | Setpoint optimization and golden-batch guidance | ML with optimization; advanced process control | Higher yield and throughput; lower variability |
| **Energy** | Load forecasting and utility optimization | Time-series forecasting, optimization | Lower energy per unit; peak demand control |
| **Planning** | Demand forecasting and production scheduling | Forecasting, optimization | Better service levels; fewer changeovers |
| **Safety** | PPE compliance and restricted-zone alerts | Computer vision | Faster response to unsafe situations |
| **Knowledge** | Assistants for procedures, manuals and troubleshooting | Generative AI with RAG | Faster answers; knowledge retained |

## 4 CHOOSING THE RIGHT USE CASES

Score each candidate on two axes: the business value of solving the problem, and how ready the data, people and process are. Start in the top-right quadrant. High-value ideas with poor readiness become strategic bets: fund the data foundation first, then return to them.

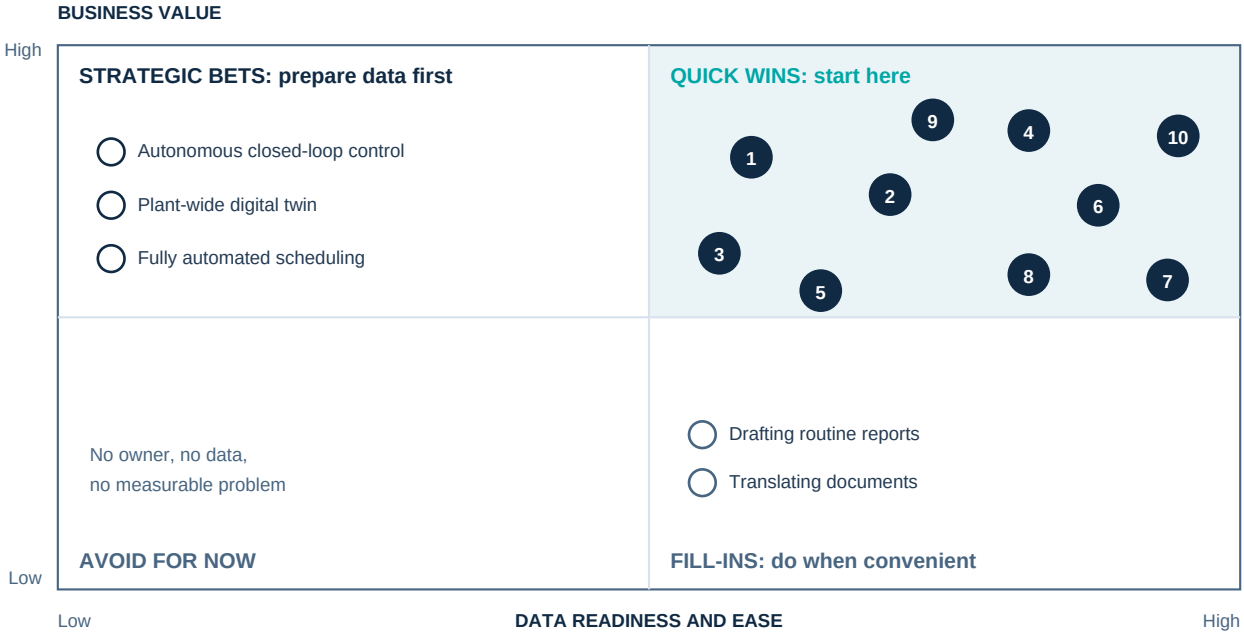


**Figure.** Prioritization matrix. Numbers refer to the quick wins in Section 6. Positions are indicative; assess value and readiness for your own site.

| Criterion | Questions to score |
|---|---|
| **Value** | How often does the decision or defect occur? What does one occurrence cost in output, quality, energy or safety? |
| **Data readiness** | Is the data available, in context and of good quality? Are there labels or outcomes to learn from? |
| **Ease of adoption** | Who will use the result, where, and how does it fit their existing work? Is an owner willing to lead? |
| **Risk** | What happens if the model is wrong? Can a person review the output before it takes effect? |

## 5 REFERENCE ARCHITECTURE FOR INDUSTRIAL AI

A layered architecture lets each new AI use case reuse the same data, tools and controls. It applies to discrete, process and hybrid plants and to on-site, cloud or hybrid deployment. A trust layer of security, AI governance, data protection and human oversight spans every layer.

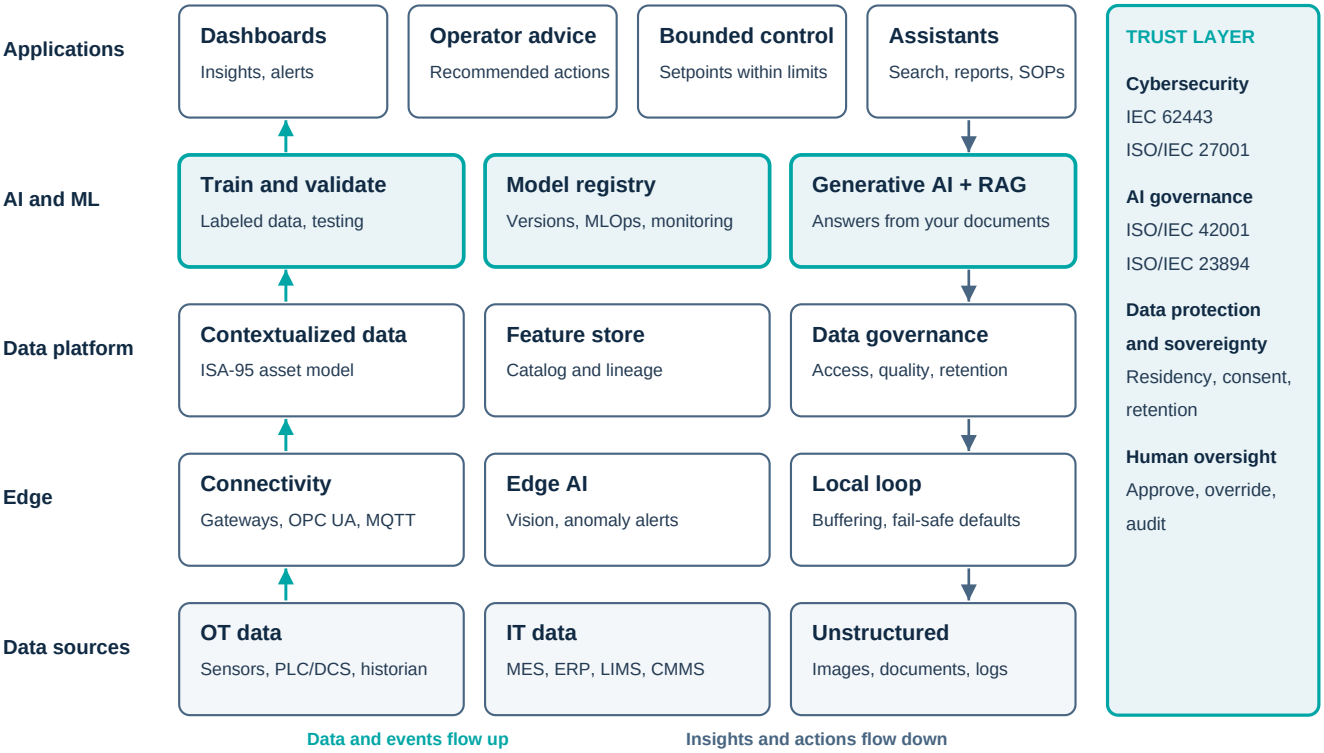


**Figure.** Reference architecture for industrial AI. Data flows up from plant and business systems; insights and bounded actions flow back down; trust controls apply at every layer.

- **Edge:** runs time-critical models (vision, anomaly alerts) near the machine and falls back to safe defaults if a model or link fails.
- **Data platform:** stores data with its context (asset, product, batch, units) using an ISA-95-based model, with catalog and lineage.
- **AI and ML:** trains, validates, versions and monitors models; generative AI answers from approved documents through RAG.
- **Applications:** deliver results where work happens: dashboards, operator advice, bounded control and assistants.
- **Interfaces:** OPC UA and MQTT for plant data, standard interfaces to MES, ERP and maintenance systems, and the Asset Administration Shell (IEC 63278) for asset information.

## 6 TEN QUICK WINS

These use cases are proven, need limited investment, and show tangible results within weeks to a few months. Most start at autonomy level 1 or 2. Benefits depend on the site; measure a baseline first.

| # | Quick win | Problem it solves | AI approach and data | Tangible outcome | Time to value |
|---|---|---|---|---|---|
| **1** | Visual inspection at one critical station | Manual checks miss defects or slow the line | Computer vision on a fixed camera; a few hundred labeled good and bad images | Every unit checked; defect images logged for root cause | 6–12 weeks |
| **2** | Anomaly alerts on critical rotating equipment | Unexpected failures of pumps, fans, motors | Unsupervised ML on vibration, temperature and current | Warning before failure; repairs planned | 4–8 weeks |
| **3** | Soft sensor for a lab-measured quality | Quality known hours after production | Regression on process data paired with lab history | Live quality estimate; earlier correction | 6–10 weeks |
| **4** | Energy forecasting for utilities | Peak charges, compressors and chillers poorly matched to demand | Time-series forecasting on meter, weather and production data | Better load scheduling; lower peaks | 4–8 weeks |
| **5** | Golden-batch or best-run guidance | Yield and cycle time vary by shift and crew | ML comparing current run with best historical runs | Operator advice to stay near best practice | 8–12 weeks |
| **6** | Generative AI assistant for procedures and manuals | Time lost searching documents; knowledge in few heads | LLM with RAG over approved SOPs and manuals, hosted on-site or in-country | Quick, cited answers at the line | 4–8 weeks |
| **7** | Maintenance work order text analytics | Free-text failure records cannot be analyzed | Language models classify failure modes and find repeat failures | Clear bad-actor list; better failure codes | 3–6 weeks |
| **8** | Alarm flood and nuisance alarm analysis | Operators overwhelmed by alarms | Pattern mining on alarm and event logs (ISA-18.2 context) | Targets for alarm rationalization | 3–6 weeks |
| **9** | Demand forecasting for production planning | Over- and under-production; frequent re-planning | Time-series ML on orders, history and seasonality | More stable plans; lower stock | 6–10 weeks |
| **10** | PPE and restricted-zone safety alerts | Unsafe acts noticed late | Computer vision with privacy by design (blurring, no identity) | Real-time alerts; trend data for safety teams | 6–10 weeks |

**Worked example: label inspection on a packaging line**

A packaging line relies on one operator to check label position and print quality. A fixed camera and lighting are added at the labeler, and a vision model is trained on a few hundred images of good and defective labels collected over two weeks. In the first phase the system only flags suspected defects (level 1). Once its decisions match the quality team's judgment, it is allowed to divert rejects automatically (level 4, bounded). Every defect image is stored, giving the team evidence to fix the labeler itself.

**Worked example: soft sensor in a continuous process**

A polymer plant measures a key quality property in the lab every four hours. A soft-sensor model is trained on two years of process data aligned with lab results and validated on a period it has never seen. It now shows an estimate every minute with a confidence band. Operators correct drift long before the next lab sample, while the lab result continues to confirm and recalibrate the model.

**Worked example: on-site knowledge assistant for maintenance**

Maintenance technicians spend time searching for procedures and past fixes. An LLM assistant is deployed on site and connected, through retrieval, only to approved SOPs, equipment manuals and closed work orders. Answers cite the source document and revision. No plant data leaves the site, and access follows existing user roles.

## 7 GENERATIVE AI IN INDUSTRIAL SETTINGS

Large language models are useful wherever plant knowledge sits in text: procedures, manuals, work orders, shift logs, specifications and standards. Used with retrieval from approved sources, they give fast, cited answers. Used carelessly, they can produce confident but wrong answers or expose confidential information.

| Proven use | How it helps |
|---|---|
| **Knowledge assistant for operators and technicians** | Answers questions from approved SOPs, manuals and past work orders, citing the source |
| **Shift handover and report drafting** | Summarizes logs, alarms and events into a draft for the supervisor to check |
| **Work order summarization and coding** | Turns free-text notes into structured failure modes and causes |
| **Training and multilingual support** | Creates training material and translates instructions for a diverse workforce |
| **Engineering support** | Drafts documentation and explains existing automation logic for review by qualified engineers |

### Risks and controls

| Risk | Control |
|---|---|
| **Wrong or invented answers (hallucination)** | Answer only from retrieved, approved documents; show citations; let users flag errors |
| **Outdated information** | Connect to document control so only current revisions are retrieved; show revision in answers |
| **Confidential data exposure** | Host on site or in-country; contract that your data is not used to train the provider's models |
| **Prompt injection and misuse** | Treat documents and inputs as untrusted; limit what the assistant can access and do; log usage |
| **Over-reliance** | Keep a human reviewer for decisions; never connect a language model directly to control actions |
| **AI agents taking actions** | Limit agents to autonomy level 3 (act with approval) with narrow, audited permissions |

## 8 DATA READINESS

Most AI projects that stall do so because of data, not algorithms. The ISO/IEC 5259 series defines data quality for analytics and ML; ISO/IEC 8183 describes the data life cycle. In practice, check the following before starting.

| Check | Question to answer |
|---|---|
| **Availability** | Is the data collected at all, at the right frequency, and kept long enough? |
| **Context** | Can each value be linked to its asset, product, batch, order and operating mode? |
| **Quality** | Are there gaps, frozen values, wrong units, unsynchronized clocks or uncalibrated sensors? |
| **Labels** | For supervised learning, are there enough trusted examples of each outcome (defects, failures, lab results)? |
| **Access** | Can data be read without disturbing control systems, and through secure, documented interfaces? |
| **Ownership** | Who owns each data set, who may use it, and for what purpose? |
| **Classification** | Is the data personal, confidential process know-how, export-controlled or public? |

# 9 DATA SOVEREIGNTY AND DATA PROTECTION

Industrial AI uses two kinds of sensitive data: **personal data** about employees, contractors and customers, and **confidential operational data** that embodies process know-how, recipes and performance. Both must be protected, and many organizations must also control **where** data is stored and processed and **who** can access it.

| Concept | Meaning |
|---|---|
| **Data sovereignty** | Data is subject to the laws of the country where it is collected or stored; the owner keeps legal and practical control over it |
| **Data residency** | The physical or geographic location where data is stored and processed |
| **Data localization** | A legal requirement that certain data stays within a country |
| **Data protection and privacy** | Lawful, fair and secure processing of personal data, with defined purposes and individuals' rights |
| **Confidentiality of know-how** | Protecting recipes, process parameters and performance data as trade secrets, including from use in training third-party models |

## Regulations and standards to consider

| Instrument | Relevance |
|---|---|
| **General Data Protection Regulation (EU GDPR)** | Personal data of individuals in the EU; lawful basis, data protection impact assessments, transfer rules |
| **UAE Federal Decree-Law No. 45 of 2021 (PDPL); DIFC and ADGM data protection laws** | Personal data processing in the UAE and its financial free zones |
| **Saudi Personal Data Protection Law (PDPL)** | Personal data processing and cross-border transfer conditions in Saudi Arabia |
| **Other national laws and sector rules** | Many countries and critical sectors restrict where data may be stored; confirm for each site |
| **EU AI Act** | Risk-based obligations for AI systems; AI used as a safety component of products such as machinery can be high-risk |
| **ISO/IEC 27001 and ISO/IEC 27701** | Information security management and its privacy extension |
| **ISO/IEC 27018** | Protection of personal data in public cloud services |
| **IEC 62443** | Security of industrial automation and control systems connected to AI platforms |

Regulations change and apply differently by country, sector and data type. Confirm obligations with legal and compliance advisors for each site.

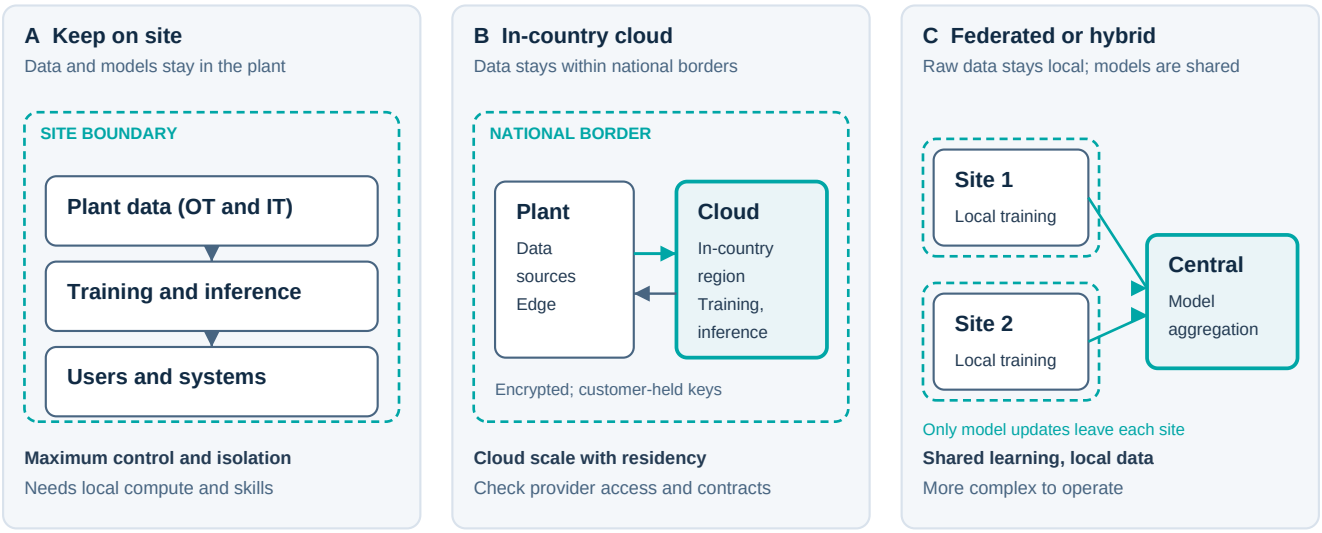


**Choose the pattern per use case and data class, not once for the whole company.**

**Figure.** Three deployment patterns that keep control of data. Many organizations combine them: sensitive process data on site, less sensitive workloads in an in-country cloud, and federated learning across sites.

### Practical controls

- **Classify data** before any project: personal, confidential know-how, export-controlled, internal or public; set rules per class.
- **Contract for control:** data location, no use of your data to train the provider's general models, ownership of trained models, export and deletion on exit.
- **Encrypt** in transit and at rest, with keys held by your organization where possible; restrict and log administrator access.
- **Minimize personal data:** blur faces, avoid identity in vision use cases, limit retention, and run a data protection impact assessment where required.
- **Govern generative AI use:** no confidential data in public AI tools; provide an approved assistant hosted on site or in-country instead.
- **Respect the OT boundary:** AI platforms read plant data through a demilitarized zone; any write-back is bounded, approved and logged.

## 10 GOVERNANCE, SAFETY AND STANDARDS

AI should be managed like any other system that affects production, quality and safety: with clear ownership, risk assessment, change control and monitoring. A growing set of international standards supports this.

| Standard or framework | Subject |
|---|---|
| **ISO/IEC 22989** | AI concepts and terminology |
| **ISO/IEC 23053** | Framework for AI systems using machine learning |
| **ISO/IEC 42001** | AI management system: policies, roles, risk, life cycle, suppliers |
| **ISO/IEC 23894** | Guidance on AI risk management |
| **ISO/IEC 5338** | AI system life cycle processes |
| **ISO/IEC 5259 series; ISO/IEC 8183** | Data quality for analytics and ML; data life cycle framework |
| **ISO/IEC 25059** | Quality model for AI systems |
| **ISO/IEC TR 24027 and TR 24028** | Bias in AI systems; trustworthiness of AI |
| **NIST AI Risk Management Framework** | Voluntary framework to map, measure, manage and govern AI risks |
| **ISO 23247; IEC 63278** | Digital twin framework for manufacturing; Asset Administration Shell |
| **ISA-95 / IEC 62264; OPC UA; MQTT** | Data models and interfaces between plant and business systems |
| **IEC 61508 / IEC 61511** | Functional safety; safety functions remain deterministic and independent of AI |
| **IEC 62443; ISO/IEC 27001** | Industrial cybersecurity; information security management |

**The safety boundary**

AI must not be part of a safety instrumented function or replace engineered protection. Where AI suggests or writes setpoints, enforce hard limits in the control system, keep operators able to override, log every action, and treat model changes as engineering changes under formal management of change.

## 11 THE AI LIFE CYCLE AND MLOPS

An AI model is not finished when it goes live. Processes, raw materials, equipment and products change, and model accuracy drifts. A disciplined life cycle keeps models useful and auditable.

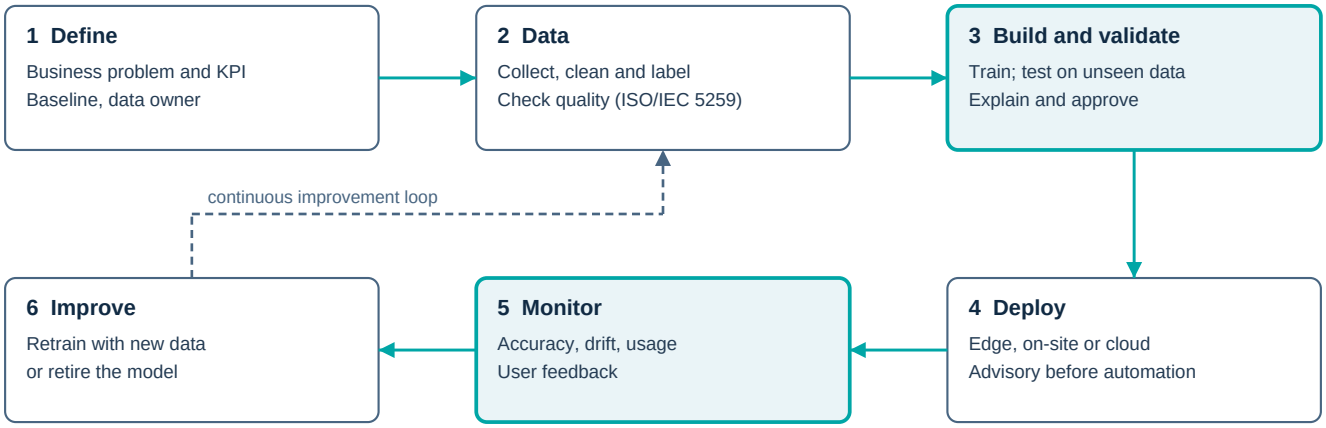


**Figure.** The AI life cycle. Monitoring triggers retraining or retirement, so models improve with use rather than decay.

- **Validate on unseen data** from a different period or line before go-live, and record the results.
- **Version everything:** data sets, code, models and configuration, so any result can be traced and reproduced.
- **Monitor in production:** accuracy against confirmed outcomes, data drift, false alarm rate and user acceptance.
- **Plan retraining** after process, product or equipment changes, and retire models that no longer add value.

## 12 ROADMAP, KPIS AND NEXT STEPS

| Phase | Activities | Output |
|---|---|---|
| **1 Identify (weeks 1–2)** | List losses and decisions; shortlist quick wins; check data readiness and classification | Two or three prioritized use cases |
| **2 Prepare (weeks 3–4)** | Agree owner, KPI and baseline; choose deployment pattern; confirm data protection and contracts | Approved pilot plan |
| **3 Pilot (weeks 5–12)** | Build and validate the model; deploy at level 1 or 2; train users; collect feedback | Model in daily use |
| **4 Prove (weeks 13–16)** | Compare with baseline; review risks; decide level of autonomy; document the template | Measured results; scale decision |
| **5 Scale** | Replicate to lines and sites; establish MLOps and governance under ISO/IEC 42001 | AI as a managed capability |

| Area | Example KPIs |
|---|---|
| **Quality** | Defect escape rate, first-pass yield, inspection coverage |
| **Availability** | Unplanned downtime, failures detected before functional failure |
| **Process** | Yield, throughput, variability of key quality variables |
| **Energy** | Energy per unit, peak demand |
| **People** | Time to find information, training time, adoption and acceptance of AI advice |
| **Model health** | Accuracy, precision and recall, false alarm rate, drift alerts |

## Pitfalls to avoid

1. Starting with a technology instead of a business problem.
2. Underestimating data preparation, context and labeling.
3. Automating before people trust the advice.
4. Ignoring data protection, residency and contracts until late.
5. Pasting confidential data into public AI tools.
6. Running pilots with no baseline or owner.
7. Deploying without monitoring for drift.
8. Letting AI touch safety functions.

### Next steps

- ☐ List the top losses where a frequent decision or inspection could be improved
- ☐ Pick two quick wins from Section 6 and check data readiness (Section 8)
- ☐ Classify the data involved and choose a deployment pattern (Section 9)
- ☐ Name an owner, a KPI and a baseline for each pilot
- ☐ Set up basic AI governance and a safety boundary before go-live

**Important:** This guide is general guidance, not legal advice or a design specification. Regulatory references are summaries; confirm requirements for each jurisdiction. For an AI readiness assessment or a quick-win pilot, contact INOVIQ Technologies at www.inoviqtech.com.