

Industrial AI: Accelerate Growth

A practical guide to building Industrial AI that connects assets, data and analytics into lasting value for modern manufacturing.

Executive summary

Manufacturers today face a balancing act. The pressure to produce more with fewer resources is relentless, while maintaining quality and reducing energy use. Rising material and utility costs continue to bite, and the pool of skilled labour is shrinking. At the same time, production equipment is becoming more sophisticated, generating ever-growing streams of data. Turning that raw information into timely, useful guidance, something operators can act on now rather than read about in a month-old report, has never been more important.

Industrial AI offers a way forward. By combining statistics and machine learning, trained on real plant data, it delivers insights where and when they are needed. It means models that predict quality before lab results land. Soft sensors that reveal what cannot be measured directly. Tools that recommend the best settings for the next run. Above all, it means providing operators and engineers with the practical support they need to make informed decisions, shift after shift, using proven methods and technologies.

In this paper, we focus on a stack built on three proven components. Kepware KEPServerEX connects to a wide range of PLCs, drives and DCSs, publishing data through open, industry-standard interfaces. HighByte Intelligence Hub adds context by modelling assets, products and batches, and

creating clean, self-describing payloads ready for use. GE Vernova Proficy CSense takes that structured data and turns it into models and recommendations. Together, they create a clear path from connectivity to context to closed-loop action.

We have chosen these solutions because they work, and we know how to maximise their benefits. With over 35 years of experience in manufacturing automation, our teams understand how to unlock value from industrial data. We take a holistic approach, examining every aspect of an operation to develop long-term programs that deliver sustained returns, rather than short-lived projects that incur additional costs over time.

We will look at two common starting points.

Greenfield sites, where the data model can be designed up front and equipment commissioned to match it. And brownfield sites, where modern tools can be layered over existing systems to lift data quality without interrupting production. For each, we outline the systems, data structures and model development approaches that keep projects on track. The goal is a practical blueprint your organisation can adapt, shaped by hard-won insights into outcomes, risks and the realities of daily operations.

“Industrial AI delivers value when it moves beyond experiments and becomes part of daily operations. The goal is not abstract models but practical tools that help teams spot problems earlier, adjust processes with confidence, and achieve results that last.”

**Cameron Bolt,
Novotek AI Expert**



Solution spotlight : Proficiency CSense

Proficiency CSense is a unified industrial analytics platform that bridges the gap between raw plant data and continuous operational improvement. It gives process engineers a single environment to analyse, monitor, predict, simulate, and optimise operations with real-time control capabilities, often described as “5-in-1” functionality.

At its core, CSense enables engineers to combine data from multiple industrial sources, sensor networks, historians, MES, and control systems, to predict outcomes and automate corrective actions. This accelerates improvements by lowering variability, boosting throughput and increasing productivity.

The platform delivers a wide range of outcomes:

- **Reducing process variability by using AI and machine learning to detect patterns and adjust processes proactively.**
- **Accelerating troubleshooting through visual analytics that quickly pinpoint issues.**
- **Minimising downtime with tools that monitor PID loops and process health.**
- **Optimising with a Process Digital Twin to simulate adjustments before applying them in production.**
- **Improving data integrity with validation and cleaning at the source.**

CSense also stands out for the speed with which solutions can be deployed. Engineers can use visual, drag-and-drop analytics, wizard-driven data mining and pre-built templates to design solutions without needing specialist programming expertise. For advanced users, integration with Python and .NET provides the flexibility to embed custom algorithms and extend functionality.

The platform is built for flexibility and scale. Deployments can be on-premises, at the edge, or in secure cloud environments. This allows both standalone plant initiatives and enterprise-wide rollouts, with a consistent architecture that adapts to organisational needs.

Industries that benefit most from CSense include food and beverage, utilities, pharmaceuticals, mining, metals and water and wastewater. In each of these sectors, it helps improve

everything from product quality and yield to energy consumption and equipment reliability.

Another strength is its ability to deliver results quickly. Many plants see useful insights within days and measurable ROI within months. Typical improvements range across key operational metrics, with gains in throughput, reduced waste and energy usage along with more stable product quality.

CSense also plays an important role in Industry 4.0 strategies. By turning time-series process data into actionable intelligence, it supports predictive maintenance, anomaly detection, loop performance monitoring, and real-time decision support. This enables organisations to move steadily from reactive operations towards proactive and, in some cases, semi-autonomous optimisation.

“Proficiency CSense is a flexible and scalable platform that unifies analytics, monitoring, prediction, simulation, and optimisation in a single package. It empowers process engineers to convert industrial data into tangible value quickly, improving quality, productivity, energy efficiency and asset reliability. With its combination of fast deployment, intuitive tools, and advanced modelling capabilities, CSense provides a practical route to embedding Industrial AI into everyday operations and achieving lasting business outcomes.”

**Dr Martin Paczona,
Head of Industrial
Data Science**



GE VERNOVA

Distributor

Proficiency Software & Services

The greenfield path

A greenfield build allows you to set the foundations with care. The aim is a guided sequence that moves from intent to stable operations without rework. The steps below are written so a project manager can run them as a plan, and so engineering leads can see exactly what to prepare and when.



1. Define outcomes and scope

- Agree the first three business results to target, such as yield uplift on a product family, reduced start-up losses, or lower energy per unit.
- Select the inaugural production area and one secondary area ready to follow.
- Capture constraints early. Note utility limits, quality compliance needs, supplier interfaces, and commissioning windows.
- Nominate a single owner for each result with authority to make trade-offs.



3. Time, identity and trust services

- Choose plant time sources and apply NTP or PTP across equipment that supports it.
- Stand up identity and access controls for servers and applications.
- Create a certificate plan for all secured endpoints, with renewal rules and storage locations.



5. Network and addressing

- Assign subnets for control, data, and management traffic. Document addressing for gateways and servers.
- Prepare firewall rules for required protocols and target systems, with named owners for each rule set.
- Plan redundant paths for critical segments and verify link failover behaviour.



2. Project charter and governance

- Form a cross-functional core team: controls, process, quality, OT networking, data engineering, operations, IT security.
- Set decision rhythms. Weekly technical stand-up, fortnightly steering check-in, monthly risk review.
- Establish a change process for configuration, models, and pipelines with clear approval steps.



4. Compute and storage layout

- Specify edge servers for connectivity and context services. Define CPU, memory, storage tiering, and virtualisation approach.
- Size a historian and file storage for raw and contextualised data. Include retention and backup rules.
- Reserve a staging environment for proving configurations and models



6. Vendor and builder requirements

- Add data and context deliverables to equipment specifications. Include tag lists with units, ranges, and descriptions.
- Require acceptance tests that prove the agreed information model at factory and site tests.
- Ask suppliers to identify any special interfaces or converters they will bring to site.



7. Asset and information model

- Build a hierarchical model that reflects the plant layout. Include units, equipment modules, control modules, and utilities.
- Create templates for repeated assets such as pumps, ovens, mixers, fillers, and packaging cells.
- Define naming rules for variables, alarms, states, and events. Publish a short style guide with examples.



9. Kepware connectivity commissioning

- Build a project structure that mirrors the plant model. Organise channels, devices, and tag groups for clarity.
- Configure the required device drivers, test handshakes, and capture diagnostic snapshots for records.
- Enable store-and-forward buffering for loss protection. Verify data continuity by simulating link interruptions.
- Create a golden test for each device type that checks data type, range, timestamp quality, and update rate.



11. Data readiness gates

- Run completeness checks against the inaugural use cases. Confirm coverage by tag, by state, and across the shift.
- Validate timestamp alignment across sources. Confirm time skew thresholds for acceptance.
- Produce a short data dictionary for the selected area. Include source, meaning, unit, and owner for each variable



8. Measurement plan and instrumentation

- Confirm which variables must be measured directly and which can be inferred later by soft sensors.
- Grade each instrument for criticality, expected accuracy, and calibration interval.
- Place additional sensors where models would otherwise lack coverage, such as flow, temperature, or vibration points.



10. HighByte context and pipelines

- Implement asset templates that wrap values with metadata such as units, limits, product codes, and batch identifiers.
- Publish to multiple targets from one template.
- Version models and pipelines. Promote from staging to production through a scripted process with review checkpoints.
- Add automated checks that validate required fields and flag unexpected nulls or out-of-range values.



12. CSense environment and model backlog

- Install the analytics environment and connect to live and historical data sources.
- Build a backlog of model candidates with process engineers. Prioritise soft sensors, predictors, and monitoring tasks that support the agreed outcomes.
- For each candidate, define input lists, target variables, training windows, and acceptance criteria.
- Prepare feature engineering notes based on process knowledge, including lags,



13. Model development workflow

- Split historical data into training, validation, and test periods that reflect seasonality and product mix.
- Train initial models and produce human-readable diagnostics that show drivers and sensitivities.
- Conduct structured model reviews with operations and quality to confirm plausibility and actionability.
- Record version, parameters, and data ranges for each accepted model in a registry.



14. Operator experience and handoff paths

- Decide how recommendations will appear during the shift. Options include existing SCADA screens, a lightweight web panel, or a dashboard beside standard work.
- Use simple language and a consistent format: what is predicted, confidence or range, suggested action, and a small number of key drivers.
- Provide an acknowledgment or feedback mechanism so actions and outcomes can be tracked.



19. Scale-out pattern

- Package Kepware projects, HighByte templates, and CSense model kits so they can be reused in the next area.
- Keep a catalogue of proven features, pipelines, and dashboards. Note prerequisites and common pitfalls.
- Plan the next two rollouts while lessons from hypercare are fresh.



20. Sustained operations and improvement

- Establish monthly checks for data quality, pipeline health, and model performance. Assign owners and due dates.
- Schedule periodic model refreshes when product mixes or equipment behaviour changes.
- Keep a small backlog of enhancement ideas sourced from operators and engineers. Deliver improvements on a regular cadence.



15. Closed-loop design where appropriate

- Identify loops or setpoints that qualify for automatic adjustment under guardrails.
- Define limits, rates of change, and safe states for each action. Document how to disable the loop quickly if needed.
- Run shadow mode first. Compare suggested changes against actual outcomes over an agreed period before enabling automatic writes.



16. Testing and acceptance

- Build a library of test cases that cover connectivity, context, analytics, and operator paths.
- Prove failure behaviour by interrupting links, forcing bad values, and simulating sensor outages.
- Run performance tests at expected peak tag counts and message rates. Capture latency and loss statistics.
- Conduct user acceptance testing with shift teams across day, swing, and night.



21. Documentation pack

- Maintain living documents for the asset model, data dictionary, interface list, firewall rules, pipeline diagrams, and model registry.
- Store test evidence, commissioning results, and user training records in a central location.
- Record every decision that affects process behaviour, with dates and approvers.



22. Procurement and licensing view

- Prepare a bill of materials for servers, storage, network gear, and any additional sensors.
- Track licence assignments and renewal dates for connectivity, context, and analytics components.
- Budget for staging environments and test equipment to keep changes safe.



17. Training and adoption

- Deliver short sessions for operators, engineers, and maintenance teams. Focus on how to interpret guidance and when to escalate.
- Provide job aids and quick reference cards at the line.
- Agree success signals for the first month in service, such as reduction in start-up deviations or faster recovery after disturbances.



18. Go-live and hypercare

- Sequence cutover by cell or unit. Announce stop-go points and clear rollback triggers.
- Staff a daily hypercare huddle with the core team during the first weeks. Review incidents, model behaviour, and operator feedback.
- Track a small set of metrics that reflect the chosen outcomes and publish them where teams can see progress.



23. Example timeline guide

- Weeks 1–2: outcomes, charter, and architecture decisions.
- Weeks 3–6: infrastructure build, model templates, and device tagging standards.
- Weeks 7–10: connectivity commissioning and initial context pipelines.
- Weeks 11–14: data readiness checks, model backlog and first builds.
- Weeks 15–18: operator experience, testing, go-live preparation.
- Weeks 19–22: hypercare and first scale-out area.

“Follow this path and a new site arrives at day one with clean data, clear context, and practical analytics in service of defined results. The programme then grows by repeating the same tested steps with less effort each time.”

**Cameron Bolt,
Novotek AI Expert**



The brownfield path

An existing plant presents different challenges from a greenfield build. The focus is on introducing modern connectivity, context, and analytics without disrupting operations or compromising legacy systems. The path below is designed to give structure to that process, moving from discovery to sustained improvement.



2. Asset and system inventory

- List all assets, including vendor, model, firmware or software version, and current communication method.
- Document current data consumers: SCADA, MES, historian, spreadsheets, third-party monitoring tools.
- Mark critical equipment with no existing data link as high-priority for connection planning.



4. Data quality baseline

- Select tags from each major system and record completeness, update rates, and timestamp alignment.
- Note patterns of missing data, irregular intervals, or incorrect units.
- Identify sensors or instruments with recurring faults or calibration gaps.



6. Introduce centralised connectivity with Kepware

- Route existing data sources through Kepware where possible to simplify interfaces.
- Enable store-and-forward for resilience and set up diagnostic logging for troubleshooting.



1. Establish project intent and boundaries

- Define the first operational outcome to target, such as stabilising a critical process, reducing downtime on a specific line, or improving quality consistency.
- Select the scope area where changes can be made safely within planned maintenance windows.



3. Network and access review

- Map the control network, data routes, firewalls, and gateways.
- Verify whether secure protocols (OPC UA with encryption, TLS) are supported natively or require an intermediary.
- Document current access controls for devices and servers.



5. Quick-win targeting

- Work with process and maintenance teams to identify a problem that is measurable, high-impact, and well understood by operators.
- Choose a case that can be addressed with improved visibility, predictive alerts, or performance monitoring without requiring plant-wide changes.



7. Layer in HighByte Intelligence Hub

- Enrich incoming data with metadata such as product, batch, and operational state.
- Publish contextualised payloads to the historian, reporting tools, and selected analytics targets.
- Maintain existing point-to-point links until the new pipelines prove stable.



8. Data validation and readiness

- Compare data from new pipelines to existing systems to confirm accuracy.
- Verify that context fields are populated consistently for all assets in scope.
- Log and correct any mismatches between physical measurements and reported values. can be made safely within planned maintenance windows.




10. Expand coverage gradually

- Add more assets in phases, prioritising areas with strong operator engagement or pressing business needs.
- Migrate consumers from legacy data feeds to the new contextualised pipelines.
- Introduce additional CSense models for monitoring, prediction, or optimisation as data coverage improves.



12. Safeguards and risk control

- Keep legacy systems operational until the new architecture is accepted by all stakeholders.
- Test failover paths so data continues flowing if new components go offline.
- Document rollback procedures for every change.



9. Deploy a focused CSense pilot

- Connect CSense to historical and live contextualised data.
- Build a model for the chosen quick-win problem, using contribution analysis and visualisation to make drivers clear to the plant team.
- Run the model in advisory mode, tracking recommendations and outcomes over a set period.
- Share results openly to build trust.



11. Integration with existing workflows

- Adapt visualisation tools or SCADA screens to present AI-driven insights alongside familiar process displays.
- Train operators and engineers on interpreting recommendations, including when to escalate or override.
- Establish feedback loops so operator observations feed back into model refinement.



13. Scaling pattern for brownfield sites

- Maintain a catalogue of configurations, templates and CSense models that worked in one area.
- Reuse these patterns, adjusting only for site-specific variables.
- Keep lessons learned in a shared repository.

“By following this structured approach, brownfield plants can progressively modernise their data infrastructure and deploy AI without risking the stability of production. The key is measured progress, clear communication with stakeholders, and early wins that demonstrate the value of change.”

Dr Martin Paczona, Head of Industrial Data Science



Security and compliance by design

Security in industrial AI projects is not just about protecting data; it is about safeguarding the operation itself. Every connection, service, and application must be assessed for potential risk and designed with mitigation in mind.

The first step is to separate operational networks from enterprise and external systems, allowing data to flow only through controlled points. These gateways should enforce encryption, authentication, and strict access permissions. Certificates must be managed and renewed as part of routine maintenance.

Compliance requirements vary by industry but often include traceability of data, audit trails for system changes, and secure storage of historical

records. Meeting these requirements is easier when security measures are designed in from the start rather than added later.

Monitoring is a vital part of security. Continuous logging of data flows, configuration changes, and access attempts enables early detection of anomalies. Regular reviews of these logs, combined with vulnerability assessments and patch schedules, ensure systems remain resilient against evolving threats.

By treating security and compliance as an integral part of the architecture rather than a separate layer, organisations can reduce risk without slowing down the adoption of industrial AI.

"Security and compliance are not just technical requirements, they are the foundation of trust in operations. When we introduce new data flows and analytics, we are effectively expanding the plant's digital perimeter. Every additional connection, every pipeline, and every endpoint must be treated as part of the security design. The approach we take is to embed these considerations at the start of any Industrial AI programme rather than treating them as an add-on."

Security begins with network segmentation and the use of secure protocols such as OPC UA with proper certificate management. We also place strong emphasis on access controls, audit trails, and encryption of data both in transit and at rest. These measures protect against cyber threats but also ensure that any data used in analysis is reliable and traceable.

Compliance is equally important, especially in regulated industries where records, auditability, and governance are non-negotiable. AI deployments must support these standards by preserving lineage, logging model changes, and documenting how data has been used to make decisions. Our role is to help customers implement systems that meet both internal policy and external regulation without slowing down innovation.

By designing for security and compliance from day one, organisations reduce risk, avoid costly rework, and build confidence in their AI-driven processes. Ultimately, it is about creating an environment where operators, engineers, and management can trust the insights being generated and act on them without hesitation."

Cameron Bolt, Novotek AI Expert



Use case cheat sheet

A quick reference to proven Industrial AI applications that show where value can be captured fastest.

Predictive quality

Models forecast key quality metrics before lab confirmation, allowing adjustments during production to prevent defects and reduce waste.



Condition monitoring

Continuous tracking of equipment health indicators enables early intervention before failure, reducing unplanned downtime.



Recipe and changeover efficiency

Analysis of sequence timing and equipment readiness reduces downtime between product runs.



Prioritise relevant alarms

Grouping, filtering, and prioritising alarms improves operator response and prevents nuisance or fatigue conditions.



Throughput and yield improvement

Monitoring process stability and identifying variability drivers supports consistent output and higher line efficiency.



Energy optimisation

Real-time analysis of utility consumption identifies inefficiencies and recommends setpoints or operational changes to minimise costs.



Control loop performance

Automated detection of oscillation, drift, or saturation helps prioritise retuning work for loops that have the highest operational impact.

In brief: AI deployment

A concise overview of the steps, patterns, and best practices that turn Industrial AI from concept into daily reality.

Edge-first

Connectivity and analytics run on-site close to equipment, reducing latency and keeping data within plant boundaries. Suitable for processes requiring immediate response or where cloud access is limited.



Hybrid

Connectivity and context are handled locally, with selected datasets sent to cloud services for model training or long-term analysis. Balances local control with centralised capabilities.



Centralised

Multiple plants feed data into a shared infrastructure for analytics, enabling standardised models and cross-site benchmarking. Works best with stable, high-bandwidth connections.



Phased adoption

New layers are added incrementally, starting with connectivity, then context, and finally analytics. Allows early value capture and reduces change risk.



Redundant architecture

Duplicate servers and failover configurations ensure continuity in the event of equipment or network issues, protecting critical data flows.

Operating model

A successful Industrial AI programme relies on clear roles, disciplined processes, and the ability to sustain improvements once systems are in place. This means combining traditional operational responsibilities with new practices for managing data pipelines and machine learning models.



Roles and responsibilities

Operational Technology (OT) manages connections to plant equipment, ensures sensor accuracy, and maintains control network stability.



Data engineering designs and maintains the context layer, keeping asset models, pipelines, and data quality checks up to date.



Process engineering works directly with analytics tools, developing and refining models to address operational challenges.



Project or product owners coordinate efforts, track progress against objectives, and make decisions on scaling or adjusting scope.



IT and security maintain infrastructure availability, enforce access controls, and ensure compliance with company and regulatory standards.



Model lifecycle management

Maintain a central registry for all models, including version, deployment date, target variables, and input features.



Document acceptance criteria before a model is deployed, covering accuracy, reliability, and operator usability.



Use a staged deployment approach: develop and test in a controlled environment, run in advisory mode for validation, and then move into production with clear rollback procedures.



Track performance over time and flag models for retraining when drift or reduced accuracy is detected.



Version-control connectivity configurations, asset templates, and context pipelines to ensure reproducibility.



Establish promotion steps from development to production, with automated validation to check schema, completeness, and data quality.



Monitor pipeline health and set alerts for lag, failure, or missing data.



Continuous improvement

Schedule regular reviews of model performance with operators and engineers to capture feedback.



Keep a backlog of potential use cases, enhancements, and data quality improvements, prioritising those with measurable business impact.



Standardise successful patterns so they can be applied consistently across lines or plants. Operational integration



Deliver model outputs in the tools operators already use, with clear instructions and minimal extra steps.



Provide a feedback mechanism so users can log issues or observations that feed into model refinement.



Recognise and share operational successes to maintain engagement and support for further adoption.



When the operating model is embedded into daily routines and MLOps practices are consistently applied, Industrial AI becomes a stable, trusted part of operations rather than a one-off project.

ROI model and business case

A credible return-on-investment model is essential to justify Industrial AI projects and secure continued support from stakeholders. It should directly connect technical outcomes to financial results and be realistic about both the benefits and costs.

Define measurable baselines

Select KPIs that link directly to business goals, such as scrap rate, unplanned downtime, throughput per hour, energy consumed per unit, or maintenance cost per month.



Measure and record current performance before making any changes, ensuring the data covers a representative operating period.



Include contextual factors, such as product mix, seasonal demand, or equipment condition, so that results can be interpreted accurately.



Estimate achievable improvements

Consult with process engineers and operators to validate assumptions about what is technically possible.



Use historical data to understand the normal variability of each KPI and set a realistic target for improvement.



Where possible, reference results from similar plants or previous projects to strengthen credibility.



Account for all costs

Include hardware, software licences, sensors, and infrastructure upgrades required for connectivity, context, and analytics.



Factor in internal time for commissioning, model development, validation, and operator training.



Present benefits and costs over a defined time frame, typically three to five years. Include any ongoing costs for cloud services, support contracts, or data storage.



Structure the business case

Show both the payback period (how long to recoup the investment) and the internal rate of return (IRR).



Include best-case, expected-case, and conservative scenarios to demonstrate awareness of uncertainty.



Track and report benefits

During pilot phases, log every action taken because of AI insights and the measurable effect of that action.



Compare pilot results to control lines or pre-project baselines.



After full deployment, maintain regular reporting against KPIs so benefits remain visible and attributable.



Sustain value over time

Reassess ROI annually, factoring in improvements from new models, expanded scope, or refined processes.



Track both cost avoidance and direct savings, such as avoided downtime or reduced overtime.



Use the results to prioritise further investment and to keep decision-makers engaged.



"A strong ROI model is not just about winning approval at the start, it becomes the reference point for measuring a project's long-term contribution to the business."

Dr Martin Paczosa,
Head of Industrial Data Science



Risks, constraints, and how to de-risk

Every Industrial AI deployment faces practical and organisational challenges. Identifying them early and incorporating mitigation into the plan reduces the likelihood of delays, overspending, or operational issues.

A well-prepared project plan will include a risk register covering these areas, with assigned owners and review points. Addressing risks systematically increases the likelihood of delivering reliable, lasting value from Industrial AI.

1. Data quality risks

- **Missing or incomplete data:** Some tags may be unavailable or unreliable due to sensor faults or lack of connectivity.

Mitigation: Conduct a data readiness assessment before model development. Replace or repair critical sensors in advance to prevent potential issues.

- **Inconsistent timestamps:** Devices on different clocks cause misalignment in datasets.

Mitigation: Apply time synchronisation across all relevant devices and servers.

- **Uncalibrated or drifting sensors:** Leads to models trained on inaccurate data.

Mitigation: Build calibration checks into maintenance routines and log calibration dates in the asset model

2. Connectivity and infrastructure risks

- **Legacy protocols with limited security:** Older equipment may not support encryption or modern authentication.

Mitigation: Use Kepware or secure gateways to translate and secure these links without altering the equipment itself.

- **Network bottlenecks or single points of failure:** Can lead to data loss or delays.

Mitigation: Review network architecture, add redundancy where needed, and enable buffering at the connectivity layer.

3. Model and analytics risks

- **Overfitting to historical conditions:** Models perform well on past data but fail with new conditions.

Mitigation: Use diverse datasets for training, reduce the model complexity, discuss with domain experts to select the tags that should be included. Monitor performance and retrain when drift occurs.

- **Lack of operator trust:** Models are ignored if outputs seem incorrect or unclear.

Mitigation: Involve operators early, explain model drivers, and run in advisory mode before enabling automation.

4. Organisational and cultural constraints

- **Resistance to change:** Staff may view AI as a threat to jobs or established workflows.

Mitigation: Communicate benefits clearly, focus on support rather than replacement, and involve staff in use-case selection.

- **Limited skills for ongoing support:** Without trained personnel, systems may degrade over time.

Mitigation: Include training and documentation as part of the initial deployment and refresh them regularly.

5. Security and compliance risks

- **Cybersecurity exposure:** New connections expand the potential attack surface.

Mitigation: Enforce access controls, use encrypted protocols, utilise edge deployment and conduct vulnerability assessments.

- **Regulatory breaches:** Mishandled data may violate industry regulations.

Mitigation: Map compliance requirements early and design the system to meet or exceed them.

6. Scope and delivery risks

- **Scope creep:** Adding features or areas without revisiting timelines or resources.

Mitigation: Lock scope for each phase and review lessons before expanding.

- **Over-reliance on vendor availability:** Delays in integration or support can stall progress.

Mitigation: Maintain internal capability for configuration and troubleshooting wherever possible.


7. Long-term sustainment risks

- **Neglected monitoring:** Pipelines and models can degrade silently if health checks are ignored.

Mitigation: Assign ownership for monitoring and reporting. Include XAI methods to ensure model explainability

- **Loss of knowledge:** Staff turnover can erode understanding of system design and operation.

Mitigation: Maintain up-to-date documentation and ensure knowledge is spread across multiple people.

A worker in a high-visibility yellow jacket and safety glasses is looking at a tablet. The image is overlaid with digital graphics, including binary code (0s and 1s) and circuit-like patterns, suggesting a focus on technology and data in an industrial setting.

Moving from insight to action

Over the course of this guide, we have looked at the practical steps for introducing Industrial AI into both greenfield and brownfield environments. We explored the foundations of connectivity, the importance of structured and contextualised data, the role of analytics in delivering measurable outcomes, and the governance required to sustain improvements over time. The case studies and use-case playbook demonstrate that the technology is proven and the benefits are tangible – from greater throughput and quality to energy savings and reduced downtime.

The conclusion is clear: Industrial AI is not a distant concept, but a practical toolset that can be deployed today. Success comes from combining reliable data capture, strong contextual modelling, and advanced analytics with the daily expertise of your operational teams.

Why Novotek

For more than 35 years, Novotek has been helping manufacturers transform operations through automation, connectivity, and data solutions. We understand the realities of complex production environments and have delivered programmes that balance innovation with practicality. Our partnerships with leading technology providers, including GE Vernova, HighByte, and PTC, mean we bring the best tools to each project, backed by deep implementation experience.

What sets us apart is our holistic approach. We do not treat AI as a one-off initiative. Instead, we build long-term programmes that evolve with your business, delivering sustainable improvements in performance, efficiency, and compliance.

Next Steps

The best way to start is with a focused conversation about your plant's priorities and constraints. Our AI Consultancy team is ready to work with you to identify high-value use cases, assess data readiness, and design pilot projects that deliver quick wins while building a foundation for long-term value.

Whether you are designing a new site or modernising existing operations, we can guide you through the first steps and provide a roadmap to scale AI across your organisation.

Talk to our AI Consultancy team today and take the first step towards turning Industrial AI into lasting results.

