The Intelligence System for Industrial Assets

Why manufacturers need a new intelligence layer for their industrial assets

Executive summary

Manufacturers have invested heavily in automation, data infrastructure and specialised software, yet engineers still have to bring information together manually to understand what is happening within their assets. Data sits across historians, production systems, maintenance applications and technical documentation, while critical operational knowledge remains with individual experts. The result is an industrial intelligence gap between the information manufacturers possess and their ability to turn it into understanding, decisions and action.

Marentis is building **The Intelligence System for Industrial Assets** to close this gap. It connects operational data, technical knowledge, specialised analytical methods and accumulated experience around a persistent understanding of the asset. Instead of leaving engineers to combine dashboards, models, documents and previous investigations themselves, Marentis brings evidence together in one workflow, keeps conclusions traceable and turns the outcome of each investigation into reusable operational knowledge.

In an active deployment with a large automotive manufacturer, investigations that previously took several days were reduced to a matter of hours, while participating production and maintenance engineers reported an overall 30 percent productivity gain in their daily work.

Marentis' mission is bringing the AI revolution to the machines that drive the physical world, empowering every engineer and operator to achieve more. **Its vision** is an intelligence layer through which assets can explain their condition, preserve what the organisation learns about them and progressively support the work required to operate and improve them.

Over time, this foundation enables a controlled progression from assisted intelligence to agentic and autonomous intelligence. The long-term goal is metacognitive industrial intelligence: a system that understands what it knows about an asset, recognises uncertainty, learns from previous decisions and outcomes, and progressively takes responsibility for defined operational tasks within explicit human, safety and governance controls.

For manufacturers evaluating Industrial AI, the relevant test is not the sophistication of an individual model or demonstration. The system should improve real operational work, fit existing industrial environments, keep conclusions traceable to evidence and build reusable asset knowledge that becomes more valuable over time.

1. The industrial intelligence gap

Manufacturing has become highly automated, but the way people understand and work with industrial equipment remains largely unchanged.

A modern plant may contain thousands of sensors, PLCs, robots, controllers, cameras and software systems. Operational data sits in historians, SCADA systems, MES platforms, quality databases and maintenance applications. Manuals, drawings and service instructions are stored elsewhere. Engineers and operators add another source of knowledge through their experience with the equipment.

When an asset behaves unexpectedly, someone still has to connect these sources. An engineer may begin with an alarm, inspect telemetry, review maintenance history, search the machine manufacturer's manual and ask whether the issue has occurred before. An engineer may also look at analytics tools built internally or by vendors, spending additional time interpreting the output, identifying a root cause and turning it into a concrete action. It is this last step that matters most. The answer to an issue may be found, but the process is slow, difficult to repeat and dependent on a small number of experienced people.

This is the industrial intelligence gap: the difference between the information a manufacturer has and its ability to use that information to understand what is happening, determine why it is happening and decide what to do next.

More data has not created a shared understanding

The underlying problem is fragmentation. Production, quality, maintenance and engineering are commonly supported by separate systems with different data structures and responsibilities. NIST continues to identify industrial data complexity, system integration and usable context as barriers to the broader use of AI in manufacturing.[1]

The same problem exists within an individual asset. A temperature value has limited meaning without knowing the operating state, product variant, process step and recent changes to the machine. An alarm becomes more useful when it is connected to the relevant component, surrounding signals, technical documentation and previous incidents.

Most manufacturers already hold much of the information needed to understand an issue. What they lack is a way to bring the relevant information together when it is needed.

Expert capacity is spent rebuilding context

Across Marentis customer conversations and deployments, manufacturers estimate that experienced engineers spend between 20 and 30 percent of their working time searching for data, reconstructing events and gathering the context needed to investigate operational issues. That is roughly one to one and a half working days per engineer each week.

This time goes into finding the right signals, comparing information across systems, locating relevant documentation and previous incidents, and assembling the context needed before analysis can begin.

When the required knowledge is not available internally, manufacturers depend on machine suppliers, system integrators and external service specialists. In Marentis customer conversations, external industrial specialists are commonly quoted at between €1,500 and €2,500 per day, before travel expenses or minimum service commitments. External support will remain necessary, but manufacturers should not need it for questions that could be answered from their own data and knowledge.

Knowledge is hard to preserve and scale

Experienced employees build detailed mental models of their equipment. They know which signals matter, which alarms are misleading and how changes in one part of a process affect another. Much of this understanding is learned through years of exposure to failures and unusual operating conditions.

It is rarely captured in a form that others can use. Findings remain in notebooks, email threads, presentations or individual memory. When an issue returns on another shift, line or site, the organisation may start again.

This matters as manufacturers struggle to recruit and retain skilled employees. Deloitte and the Manufacturing Institute estimate that US manufacturing may need 3.8 million new workers between 2024 and 2033, with up to 1.9 million positions potentially left unfilled.[2] Hiring and training remain necessary, but manufacturers also need to make existing knowledge easier to access.

When an operator or engineer notices a problem, the organisation should be able to answer three questions quickly: What is happening? Why is it happening? What should we do next? Today, those answers often exist somewhere inside the company. The challenge is finding them, connecting them and applying them without consuming a large share of engineering capacity. Over time, the same connected understanding can support systems that take over defined operational tasks under clear human control.

2. Why existing approaches fall short

Closing this gap does not start from nothing. Plants already use monitoring systems, historians, data platforms, analytics tools, digital twins, maintenance software and, increasingly, generative AI applications.

Each solves a valid problem. Most, however, address only one part of this gap. The engineer remains responsible for connecting the outputs and reconstructing the operational context before a reliable decision can be made.

Visibility and access do not create context

Dashboards and monitoring systems make production data visible. They show alarms, trends, process values and performance indicators and are often the first place an operator or engineer looks when something changes.

A dashboard may show that pressure increased, cycle time deteriorated and a quality threshold was exceeded. It rarely establishes whether these events are connected, which earlier change may have caused them or whether a similar pattern occurred before. The user still has to identify the relevant period and bring in maintenance records, manuals and information from other systems.

Historians, data lakes and industrial data platforms improve access to large volumes of production data. They are important foundations for analytics and AI. But collecting information in one infrastructure does not create a shared understanding of the equipment that generated it. A tag name or database field does not fully explain what a signal represents, how the component relates to the machine or which operating conditions should be considered normal.

This matters as manufacturers try to scale AI. In a 2025 survey of more than 100 manufacturing COOs, 46 percent reported limitations in their data or IT/OT systems. A quarter struggled to build reusable and scalable applications, while only 2 percent said AI was fully embedded across operations.[3]

Specialised models remain isolated

Manufacturers use specialised models for anomaly detection, predictive maintenance, quality and process optimisation. These models can create value when the problem is clearly defined and suitable data is available.

Their main limitation is isolation from the wider asset context. An anomaly model may detect an unusual signal pattern without access to machine structure, maintenance history or relevant documentation. A quality model may identify relevant variables without knowing what an operator adjusted during the shift.

These models are expensive to maintain and can also require ongoing fine-tuning or retraining as the underlying data and processes change.

Each model contributes a useful capability, but its output often becomes another piece of evidence for an engineer to interpret. It does not know when another method would be more suitable or how to combine its findings with the wider operational context.

Search and generic AI lack asset context

Modern retrieval systems and large language models make manuals, work instructions, service reports and maintenance records easier to search, summarise and discuss. But document access alone does not reveal which signals changed before an alarm, what maintenance was performed recently or whether the machine is operating outside its usual regime. It improves retrieval without creating a reliable view of live machine behaviour or the relationships between components, signals and operating states.

From isolated tools to a connected system

Dashboards, historians, analytics models, document systems and language models will remain important parts of industrial technology. The problem is that they are usually deployed as separate capabilities, each interpreting signals, documents and history on its own terms. Without a shared context connecting them, the user remains the integration layer between them.

Closing this gap requires these capabilities to work together around a persistent understanding of the asset. The system must identify relevant information, connect evidence across sources and select the appropriate method for the question being addressed.

That is the difference between adding another tool and creating an intelligence system for industrial assets.

3. Why now: The five-layer AI economy

In 2026, AI is being built and deployed as a global technology stack. NVIDIA describes this stack in five layers: energy, chips, infrastructure, models and applications. Each layer depends on those below it, while practical and economic value is created when applications turn the combined stack into useful outcomes.[4]

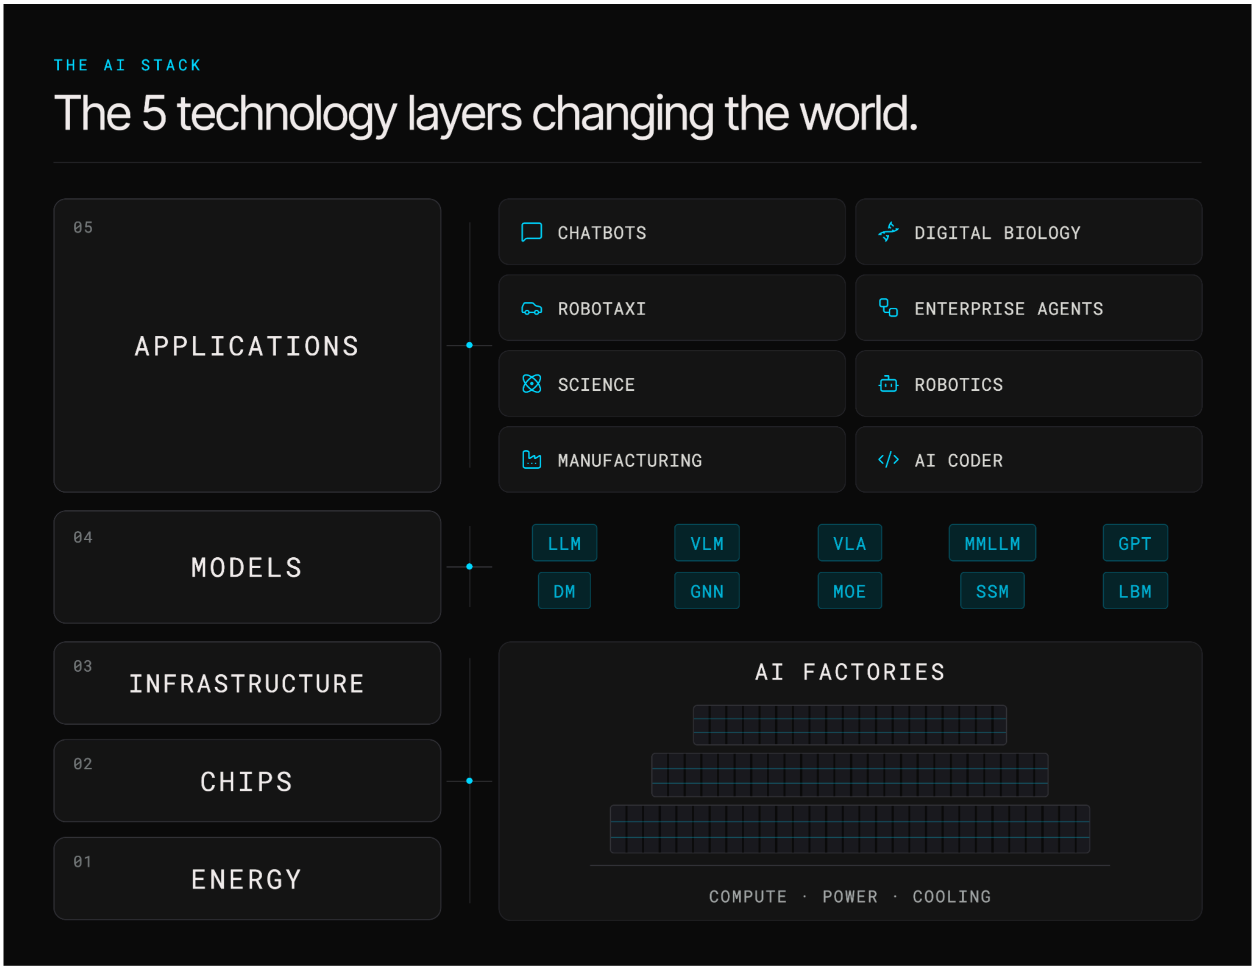

Figure 1. Illustration inspired by NVIDIA's "AI Is a 5-Layer Cake."[4]

Progress is now occurring across all five layers at the same time.

Progress across the stack

Energy powers the chips that convert electricity into computation. Chips are combined with networking, storage, cooling and software to create the infrastructure on which modern AI systems run.

Progress in these layers increases the performance, availability and cost efficiency of AI. It also expands the range of deployment options available to industrial companies, from public cloud infrastructure to private environments and systems running close to the equipment.

The model layer has expanded well beyond large language models. NVIDIA points to models that understand language, biology, chemistry, physics, robotics and the physical world. For manufacturing, this matters because documents, machine telemetry, images, process relationships and equipment behaviour cannot all be understood equally well through one model or analytical method.

At the top of the stack are applications. This is where the capabilities of the layers below are combined around real problems and where economic value is created. NVIDIA argues that models crossed an important threshold during the past year, with stronger reasoning and grounding, fewer hallucinations and applications beginning to generate measurable economic value. NVIDIA identifies manufacturing as one of the sectors where AI applications are beginning to generate practical value.[4]

Why this changes the industrial intelligence gap

The technologies described in Section 2 can now be combined more effectively. More capable infrastructure makes large-scale processing increasingly accessible. A broader range of models can interpret different types of industrial information. The application layer can connect these capabilities and apply them according to the context of the problem being addressed.

This gap has persisted because information, computing capabilities and analytical methods remained fragmented. Progress across the stack makes it possible to combine existing industrial data and knowledge in ways that were previously too limited, costly or difficult to maintain. The technical foundations for a different approach are in place.

4. Why Industrial AI requires more than an LLM

Large language models have made AI accessible through natural interaction and demonstrated how effectively models can work with technical language, documents and complex questions. In industrial environments, the central challenge is establishing an answer from the available data, technical context and operational evidence, then communicating it clearly.

Industrial intelligence is, therefore, a system problem. It requires different sources of information and analytical capabilities to work together around an asset.

Context and evidence must be part of the system

Industrial data has meaning only in relation to the equipment and the conditions under which it was generated. The same temperature, pressure or cycle time may be normal in one operating state and unusual in another. An alarm may identify where a problem was detected without revealing where it originated.

This context is specific to the industry, manufacturer, production line, team and even the role of the person asking. It also includes the informal rules, exceptions and preferences engineers rely on day to day. A useful system must therefore understand what a signal represents, which component produced it, how it relates to the wider machine or process, and the operating regime, recent maintenance and sequence of events around it.

Its conclusions must remain connected to the underlying evidence. The system should identify the signals, documents, events and previous cases that support a finding and distinguish between observed facts, analytical results and hypotheses that still need to be tested. This may require querying a historian, retrieving the correct version of a manual, comparing operating periods or applying a specialised model before an explanation can be produced.

Industrial reasoning requires different methods

Many industrial problems involve sequence, comparison and causality. What changed before a failure? Did an adjustment improve the process or merely coincide with an improvement? Does the same pattern occur across different products, shifts or operating regimes?

Answering these questions may require time-series analysis, anomaly detection, statistical comparison, graph-based methods, simulation or other specialised approaches. Possible explanations must be tested against the underlying data and the known behaviour of the asset.

Industrial reasoning is also constrained by physical reality and operational rules. A proposed action may be technically possible but unsafe under the current conditions. Another may require approval, a planned shutdown or a specific sequence of steps. NIST continues to identify reliability, explainability and the complexity of heterogeneous industrial systems as barriers to the wider use of AI in manufacturing.[1]

Knowledge must persist and remain governed

Industrial investigations create knowledge that should remain available after the immediate issue has been resolved. What was observed, which hypotheses were tested, what action was taken and whether it worked all become relevant when a similar event occurs again.

A useful industrial system should retain this history and connect new events to previous cases. Access to data, analytical tools and operational systems should follow defined permissions. Conclusions and actions should remain traceable to the evidence, people and rules involved.

From a model to an industrial intelligence system

An LLM can provide interaction, interpret technical information and communicate findings. Without connected context, however, it can also produce plausible-sounding but incorrect answers, which is why the wider system must establish those findings by connecting industrial data, asset context, specialised models, analytical tools and operational history rather than relying on the model's own responses alone.

Industrial intelligence, therefore, requires more than an LLM. It requires a system that can select the appropriate method, combine evidence across sources and maintain a reliable understanding of the asset over time.

5. Building the industrial intelligence system

These requirements point to a distinct type of industrial software: a persistent intelligence layer around the asset that connects operational information, technical context, analytical capabilities and accumulated knowledge. It reduces the work required to move from observation to a reliable conclusion by identifying relevant evidence, applying suitable methods and presenting the outcome within the asset context, moving the work toward a decision or action rather than another set of data points to interpret.

A shared understanding of the asset

The system begins with a structured representation of the equipment and its operating environment. This includes the physical structure of the asset, relationships between components, available signals, process steps, operating states and product variants. Documents, maintenance events, configuration changes and previous incidents are connected to the relevant parts of that structure.

This context persists as the equipment changes. When a component is replaced, software is updated or operating parameters are adjusted, the system can reflect the new state while retaining the earlier history.

The result is a shared reference point for data, documents, models and users. A signal is no longer only a tag in a historian. It is connected to the component that produced it, the process in which it is used and the events that may have influenced its behaviour.

Evidence and analytical capabilities in one workflow

The system must be able to access the information required for a specific question. Depending on the situation, this may include telemetry, alarms, quality measurements, maintenance records, technical documentation or earlier investigations.

Not every source needs to be copied into one central database. The system can retrieve information from existing environments under the appropriate permissions, while preserving its source and technical meaning.

It can then apply the analytical method suited to the task, drawing on the document retrieval, statistical and graph-based techniques described above, or another specialised model.

These capabilities should operate within the same investigation. Their outputs are connected to the same asset context and contribute to one evidence-based conclusion, rather than becoming separate results for the engineer to combine.

Orchestration manages the investigation

Industrial questions vary in complexity. Some require a single document or signal. Others involve several steps across different systems and analytical methods.

The orchestration layer manages this process. It determines which information and tools are relevant, controls the sequence of work and keeps each step tied to the original question.

For example, an investigation may begin by identifying the time at which behaviour changed, comparing that period with normal operation, reviewing recent maintenance and then testing possible explanations against the evidence.

The system should also recognise when the available information is insufficient. It can identify missing data, present competing hypotheses or ask an engineer for input rather than forcing a conclusion.

Operational memory makes knowledge reusable

Each investigation creates knowledge that may be useful later. The system should retain the question, evidence considered, hypotheses tested, action taken and resulting outcome.

This creates operational memory around the asset. When a similar event occurs, the system can compare it with earlier cases instead of beginning from zero. Knowledge developed by one engineer or shift becomes available to others who work with the same or comparable equipment.

Engineer feedback is part of this memory. Findings can be confirmed or rejected, and the outcome of a recommendation can be recorded. This allows the system to improve while preserving a traceable history of how conclusions were reached.

One governed interface for industrial work

For the user, these capabilities should be available through one consistent interface. Operators and engineers can ask questions in the language of their work, inspect the supporting evidence and continue the investigation where necessary.

An industrial intelligence system, therefore, combines asset understanding, evidence access, specialised methods, orchestration and operational memory within one governed environment.

Once this foundation is established, the same system can move beyond supporting investigations and begin taking responsibility for defined operational tasks.

6. A gradual path to operational autonomy

An industrial intelligence system does not need to move directly from analysis to autonomous control. Responsibility can increase gradually as the system develops stronger evidence, operational knowledge and governance. The appropriate level depends on the task, the available evidence and the consequences of being wrong.

Faster and deeper understanding

At the first level, the system helps operators, engineers and managers understand industrial situations faster and in greater depth. It can retrieve relevant information, compare operating periods, identify patterns and explain the evidence supporting a finding. The user remains responsible for interpreting the result and deciding what to do.

Decision preparation and routine work

The system can take over defined analytical and administrative work, such as preparing reports, documenting investigations, comparing alternatives or assembling evidence for maintenance and operational decisions. Accountable decisions remain with people, but the work required to prepare them becomes more consistent and efficient.

Guided execution

As confidence and operational context improve, the system can recommend next steps and support approved workflows. Human approval remains part of the process before actions affect equipment, product quality or production conditions.

Controlled autonomous action

For suitable tasks, the system can carry out defined actions within explicit permissions, operational rules and confidence thresholds. The level of human involvement should reflect the consequence of the action, with every step remaining traceable to the evidence, rules and authorisation behind it. The objective is not maximum autonomy, but the appropriate level of responsibility for each industrial task.

7. From concept to implementation

Marentis is building the industrial intelligence system described in the previous sections. It connects industrial data, technical knowledge and specialised analytical capabilities within a shared asset context that develops over time.

The system is designed to create value before autonomous action is required. Its first role is to help engineers understand industrial situations faster and in greater depth, reduce repetitive analytical work and prepare better-supported decisions.

How Marentis implements the system

Marentis connects machine data, alarms, technical documentation, maintenance history and earlier investigations within one asset context, represented as a continuously growing knowledge graph of the equipment, its relationships and the rules, exceptions and preferences engineers rely on. Its orchestration layer identifies relevant information, applies suitable analytical methods and keeps findings connected to their underlying evidence.

Engineers can begin with an operational question rather than a predefined dashboard or analysis. The system can retrieve the relevant signals and time periods, compare operating conditions, identify patterns and test possible explanations.

The investigation and its outcome become part of this expert memory, extending the knowledge graph so that previous findings can be reused and made available beyond the engineer or shift where they were first developed.

It can also support routine work such as reporting, incident documentation and preparation of evidence for maintenance or investment decisions.

Deployment, data residency and integration

Marentis is packaged as a portable, containerised deployment unit and runs within the customer's own cloud environment. This allows the deployment to follow the manufacturer's existing security, data residency and infrastructure requirements. The LLM endpoints required by Marentis can likewise be provided and fully controlled by the customer.

Operational data remains within the customer's cloud environment and is accessed through governed connections to existing data sources. Marentis can connect documents, telemetry, alarms and events, control-system records, time-series stores and enterprise applications through industrial-protocol brokers, APIs and existing data-platform connections while preserving their source and technical meaning.

The asset context and Expert Memory are customer-specific, inspectable and managed within the customer's environment. Analytical capabilities and customer-provided language models are accessed through governed interfaces, allowing manufacturers to retain control over the models, infrastructure and data used by the system rather than introducing an additional external AI environment.

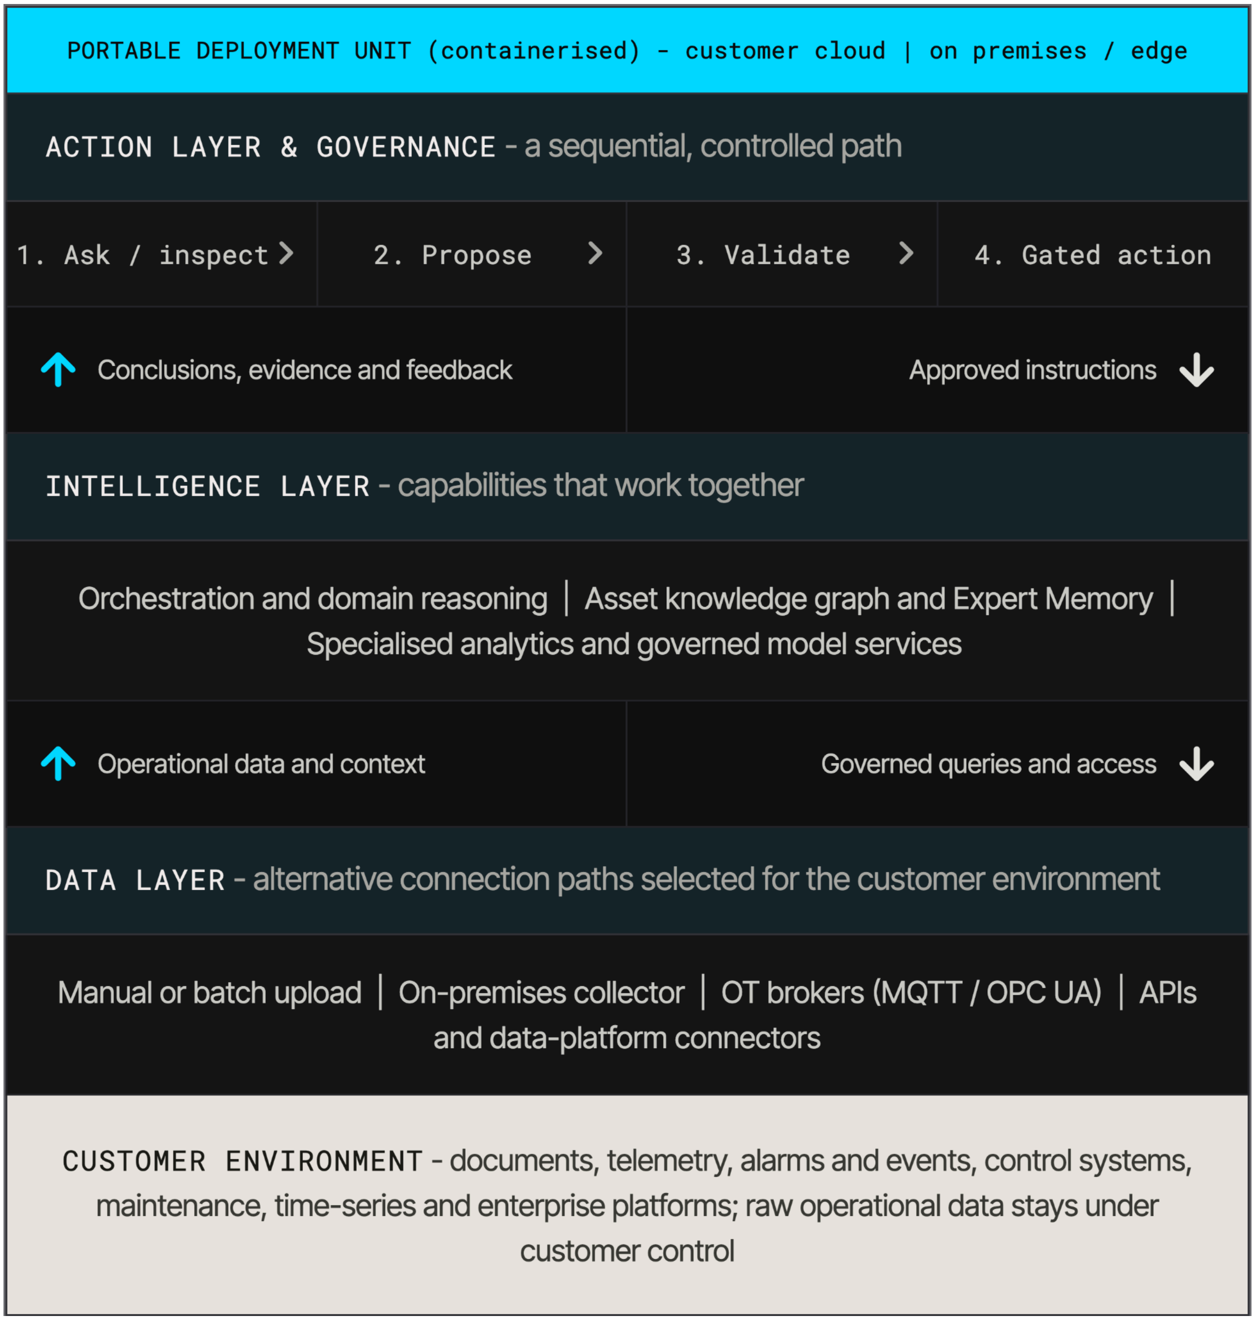

Figure 2. High-level Marentis deployment architecture

Governance across the IT/OT boundary

Connections are scoped by source, role and purpose, and deployments are designed to operate alongside existing IT/OT security policies. Access to data, tools and operational systems can be permissioned, logged and audited; conclusions retain links to the evidence, analytical steps and approvals behind them.

Human control is matched to consequence. Advice and investigation remain interactive; proposed actions can be tested against rules or, where available, a digital twin; and changes affecting equipment or production are released only through an explicit gated path. This supports greater autonomy over time without removing accountable people from consequential industrial decisions.

Where Marentis differs

Marentis' differentiation is not a single model or interface. It is the combination of persistent asset understanding, an evidence-based resolution workflow and an industrial deployment model designed for control and reuse.

DIFFERENTIATION DIMENSION	COMMON APPROACH	MARENTIS APPROACH
01 Persistent, inspectable asset understanding	Dashboards, isolated models and generic assistants expose data or answer questions, but context and knowledge often remain fragmented across people and systems.	A customer specific asset graph and Expert Memory connect equipment structure, operational rules, evidence, previous cases and outcomes. This knowledge remains inspectable, correctable and reusable as the asset evolves.
02 From evidence to governed resolution	Predictive tools often stop at an alert, while data platforms and generic agents leave the user to combine information and determine what to do next.	Marentis orchestrates telemetry, documents, operational history and specialised analytical methods around one investigation. Findings remain traceable while the system supports the path from observation to diagnosis, decision and governed action.
03 Industrial deployment and compounding value	Many approaches depend on external AI environments, continuing data science effort or integrations that remain specific to individual use cases. Knowledge created during a pilot is often difficult to reuse.	Marentis runs within the customer's cloud environment and can use customer provided LLM endpoints. Connections, asset context and Expert Memory form a reusable foundation that becomes more valuable as additional assets, workflows and operational knowledge are added.

Figure 3. Marentis differentiation from common approaches to Industrial AI and Predictive Maintenance tools

Together, these capabilities make Marentis an intelligence and resolution layer on top of existing industrial infrastructure, rather than another isolated model, dashboard or generic agent framework.

Evidence from an automotive deployment

In an active deployment with a large automotive manufacturer, Marentis was used to support team members in production and maintenance engineering. The environment already contained extensive machine data, alarms and technical knowledge, but engineers still had to identify relevant periods, inspect signals, compare operating conditions and prepare data manually before analysis could begin.

Marentis connected the relevant operational data with the asset context in one investigation workflow, allowing engineers to move from an operational question to the relevant signals, comparisons and supporting evidence without rebuilding that context manually.

Investigations that previously took several days were reduced to a matter of hours.

A recurring quality issue that once required pulling data from several systems and consulting multiple specialists could be traced to its root cause within a single shift, allowing the team to correct it before it affected further output. Participating engineers reported an overall 30 percent productivity gain across their daily work when using Marentis, and the faster diagnosis translated directly into less downtime and fewer defective parts reaching the next production step.

The immediate value came from reducing repetitive data collection, shortening the path from observation to a supported conclusion and enabling deeper investigation. Findings could be documented more consistently, while engineers gained more time for decisions requiring their experience. The deployment showed that measurable value begins with better understanding and decision preparation, before autonomous action.

From assisted to autonomous intelligence

Building on this real-world deployment experience and the graduated model of responsibility described above, Marentis is progressing along a controlled path from assisted intelligence to agentic and, ultimately, autonomous intelligence. Each stage is deployment-led: it must create value in its own right while establishing the knowledge, evidence and governance required for the next.

2026 - Industrial memory and analytical foundation. Marentis is advancing Expert Memory from customer-managed, per-user knowledge to supervised shared knowledge and cause-and-effect learning. In parallel, it is developing self-configuring, regime-aware predictive analytics for anomaly, quality and efficiency use cases, designed for customer-controlled and edge deployment. Together, these capabilities turn investigations, feedback, actions and outcomes into reusable, inspectable knowledge.

2027 - Specialized intelligence and transferable learning. Directional work includes smaller, fine-tuned specialized models, privacy-preserving knowledge abstraction, graph foundation models for industrial knowledge and expanded cause-and-effect reasoning. These capabilities are intended to deepen asset-specific intelligence while allowing useful learning to transfer across assets and deployments without transferring customer operational data between environments.

Longer term - Governed metacognitive and autonomous intelligence. As the system learns to assess its own knowledge, uncertainty and performance, it can move from answering questions to proposing and carrying out constrained actions across assets, lines and fleets. Autonomy expands only within explicit permissions, validation and human accountability.

8. What manufacturers should look for

Manufacturers evaluating Industrial AI should look beyond the quality of a demonstration or the capabilities of an individual model. The central question is whether the system can improve real operational work within the existing industrial environment.

A useful evaluation should consider how the system establishes context, supports decisions, integrates with current infrastructure and creates a foundation that can expand beyond one use case.

Start with the work, not the model

The starting point should be a recurring task that consumes meaningful engineering or operational capacity. Examples include investigating production issues, preparing maintenance decisions, documenting incidents or assembling evidence for process improvements.

The use case should have a clear baseline. Manufacturers should understand how long the work takes today, which people and systems are involved and where delays or quality problems occur.

Early success should be measured through operational outcomes such as time saved, depth of analysis, consistency of reporting and the ability to reuse earlier findings. Model accuracy alone does not show whether the system improves the work.

Require context and traceable evidence

A system is only useful if its conclusions can be trusted. Manufacturers should check that findings are traceable to the asset, operating conditions and specific evidence involved, and that observed facts are clearly distinguished from analytical results and untested hypotheses, particularly when findings justify maintenance actions, operational changes or investments.

Fit the existing industrial environment

An industrial intelligence system should work with the manufacturer's existing systems rather than require a complete replacement of them.

It should connect to historians, production systems, data platforms, document environments and maintenance applications under defined permissions. Deployment should fit within the manufacturer's controlled cloud environment while supporting customer-controlled edge components where connectivity, preprocessing or local analytical workloads require them.

Data security, access control and governance deserve particular attention. Manufacturers should expect clear control over where their data is stored and processed, which models are used, who can access which information and how that access is logged and audited. The system should integrate with existing IT/OT security policies and allow language-model endpoints to remain under customer control.

Build for reuse and controlled expansion

A pilot should create more than a successful demonstration. The connections, asset context and knowledge developed during the first use case should remain useful as the system expands to other tasks, equipment and sites.

Manufacturers should examine how quickly new data sources and assets can be added, whether previous investigations can become reusable operational memory and whether analytical capabilities can be applied across workflows. Greater responsibility should be expanded through explicit rules and approval where actions affect equipment or production.

The right system should therefore create immediate value in a defined workflow while establishing a reusable foundation for broader industrial intelligence.

9. Conclusion and the Marentis perspective

Manufacturers already possess much of the data and technical knowledge required to understand their equipment. The remaining challenge is turning that information into a reliable understanding without requiring engineers to search several systems and reconstruct the context manually.

An industrial intelligence system provides the missing layer around the asset. It connects operational data, technical documentation, analytical methods and accumulated knowledge within one governed environment. It supports the path from observation to evidence, from evidence to conclusion and from conclusion to action.

Manufacturers should begin with the work that consumes engineering capacity today. Faster investigations, deeper analysis, consistent reporting and better decision preparation can create immediate value without requiring autonomous control of equipment. The goal should be a reusable intelligence foundation that becomes more useful as it connects additional assets, workflows and operational knowledge, with responsibility increasing only as the evidence, permissions and operational controls justify it.

Marentis is building **The Intelligence System for Industrial Assets**. Its purpose is to give manufacturers a persistent understanding of their equipment and the ability to apply that understanding across operational work.

The long-term opportunity is a manufacturing environment where industrial knowledge is no longer fragmented across systems and individuals. It is available where decisions are made and increasingly supports the work required to operate, maintain and improve industrial assets. That direction is not simply toward more autonomous software, but toward metacognitive industrial intelligence that understands what it knows, learns from operational outcomes and progressively governs defined asset operations within explicit human and technical limits.

Manufacturers ready to explore what a persistent understanding of their own assets could look like are welcome to start that conversation with Marentis.

References

1. Lee, J. et al., "2026 Roadmap on Artificial Intelligence and Machine Learning for Smart Manufacturing," *Machine Learning: Engineering*, IOP Publishing, 2026.

2. Deloitte and The Manufacturing Institute, "Taking Charge: Manufacturers Support Growth with Active Workforce Strategies," April 3, 2024.

3. Cocoual, C., de Kroon, D., Lefort, F., and Chatelet, L., "From Pilots to Performance: How COOs Can Scale AI in Manufacturing," McKinsey & Company, December 15, 2025.

4. Huang, J., "AI Is a 5-Layer Cake," NVIDIA Blog, March 10, 2026.