SymphonyAI

WHITE PAPER

From Data to Decisions:

How Agentic AI with IRIS Flows is transforming industrial operations

INDUSTRIAL

From Data to Decisions: How Agentic AI with IRIS Flows is transforming industrial operations

About SymphonyAI

SymphonyAI is building the leading enterprise AI SaaS company for digital transformation across the most critical and resilient growth verticals, including retail, consumer packaged goods, finance, manufacturing, media, and IT/enterprise service management. SymphonyAI verticals have many leading enterprises as clients. Since its founding in 2017, SymphonyAI has grown rapidly to 3,000 talented leaders, data scientists, and other professionals. SymphonyAI is an SAIGroup company, backed by a $1 billion commitment from Dr. Romesh Wadhwani, a successful entrepreneur and philanthropist. More at www.symphonyai.com.

Table of contents

Introduction

Industrial operations are undergoing a profound and necessary shift. While the industry has spent decades collecting more data, deploying more sensors, and investing in increasingly sophisticated automation systems, these technologies have not solved the most pressing operational challenge: how to transform data into timely, high-quality decisions and actions that improve performance at scale.

Modern plants are more instrumented, more digitized, and more complex than ever before, yet they remain limited by systems that are static, disconnected, or overly reliant on human interpretation.

Organizations have reached the point where incremental improvements in analytics and dashboards are no longer enough.

The next leap forward requires intelligence that can interpret context, understand operational objectives, collaborate with human teams, and take action autonomously.

This new frontier of industrial innovation is agentic AI—a class of intelligent, goal-driven systems capable of perceiving, reasoning, and acting in the physical world with a high degree of autonomy and reliability.

This guide introduces IRIS Flows, SymphonyAI's agentic AI environment designed specifically for industrial operations.

IRIS Flows is much more than a workflow engine or automation tool. It is a next-generation system that enables organizations to build, orchestrate, and govern autonomous industrial agents—digital teammates that understand industrial data, collaborate across workflows, and execute decisions with precision and speed.

The Shift Toward Agentic Intelligence

For years, industrial companies have relied on dashboards, alerts, predictive models, and human expertise to identify and address operational risks. But human teams are increasingly overloaded, and predictive models often operate in isolation from real-world workflows. The consequence is predictable: insights do not always translate into consistent action, maintenance is still reactive, and knowledge is difficult to scale.

Agentic AI changes this by moving from insight generation to autonomous action. Instead of simply notifying operators about a pressure deviation or a failing pump, an agent in IRIS Flows can:

- Detect the anomaly
- Correlate it with historical maintenance logs, P&IDs, and operational thresholds
- Identify the most likely cause
- Recommend or trigger a specific maintenance workflow
- Deliver step-by-step instructions to field workers
- Confirm whether the issue was resolved
- Learn from the outcome to improve future decisions

What once required multiple specialists, tools, and manual coordination can now be executed seamlessly by a coordinated team of agents.

Why Legacy Systems Fall Short

Modern industrial operations are orchestrated across dozens of systems—historians, SCADA, PLCs, MES, CMMS, EAM, ERP, operational logs, documents, sensors, and frontline workflows. Each system sees only part of the truth. Humans are left to interpret these fragments and make decisions under time pressure. Even the most advanced predictive models depend on clean, contextualized data and consistent follow-through, both of which remain difficult to achieve.

The result is a costly gap between data and action:

- Alarms are triggered without prioritization
- Predictive warnings are missed or misinterpreted
- Maintenance instructions vary by operator
- Shift handovers lose critical information
- KPIs are influenced by hidden upstream or downstream factors
- Safety risks emerge from incomplete situational awareness

Organizations need systems that understand the operational reality end-to-end—not just isolated signals.

A New Industrial Frontier

The rise of agentic AI represents the most important industrial transformation since the introduction of automated control systems.

As organizations strive for higher reliability, agility, and resilience, IRIS Flows provides the foundation needed to move beyond traditional automation and toward true autonomous industrial operations.

This whitepaper explains the evolution that has led to this point, the technologies that enable agentic intelligence, and the unique architecture of IRIS Flows that positions SymphonyAI at the forefront of this revolution.

It explores how industrial agents work, how they collaborate, the value they create, and the future they are shaping.

The era of industrial agents has begun—and IRIS Flows is the platform that will define it.

2. The state of industrial operations today

Industrial organizations today sit at the center of converging pressures that have fundamentally reshaped the expectations placed on operations teams, technology systems, and plant leadership. From the outside, industrial facilities appear more technologically advanced than ever: machines are saturated with sensors, historians collect billions of time-series points daily, digital work instructions are replacing paper procedures, and enterprises have invested heavily in cloud migration and analytics platforms.

Yet beneath this layer of modernization lies a stark truth: many organizations still struggle to convert data into reliable, repeatable operational decisions and actions. Despite significant investment, the gap between data, insight, and execution persists. Most plants operate with a blend of automated systems, siloed data, tribal knowledge, and reactive workflows that limit their ability to perform consistently under pressure.

To understand why agentic AI represents a generational leap forward, we must first examine the realities and constraints that define industrial operations today.

2.1 Fragmented Systems and Siloed Data

Industrial environments are built on decades of layered systems—some modern, some legacy—and each holds a different view of operational truth. Time-series data lives in historians and SCADA systems. High-value engineering knowledge lives in P&IDs, 3D CAD models, and decades-old documentation. Maintenance records and failure histories reside in CMMS or EAM systems. Production execution happens inside MES systems. Workforce data lives in ERP platforms. And the most critical context often resides in the heads of long-tenured subject-matter experts.

The result is a fractured ecosystem where:

- Data lacks interoperability
- Context is inconsistent or missing entirely
- Decision-making remains dependent on manual investigation
- Automation stops at the boundary of each system

This fragmentation makes it incredibly difficult to build intelligent systems capable of understanding the operational environment holistically.

AI models trained on isolated datasets cannot account for the complex relationships between processes, assets, people, and historical patterns. To unlock true intelligence, industrial operations require contextualization—the ability to unify, relate, and interpret data across all systems, modalities, and time horizons.

2.2 Workforce Shortages and Knowledge Loss

Industrial industries face a global skills crisis. Large segments of the experienced workforce are retiring, and younger workers entering the field often lack exposure to the deep process knowledge accumulated over decades. The result is a widening gap between what operations require and the expertise available to execute safely and efficiently. Organizations increasingly report:

- Difficulty hiring and retaining qualified technicians and engineers
- Extended training timelines for new personnel
- Higher reliance on tribal knowledge
- Variability in how procedures are interpreted and executed
- Increased operational risk due to inconsistent decision-making

This erosion of institutional knowledge isn't just a workforce issue—it directly impacts asset reliability, safety, and production performance. A shrinking expert pool means plants cannot rely solely on human judgment to analyze complex patterns, recall historical lessons, or interpret nuanced signals hidden within data.

This is a core reason why agentic AI is urgently needed: to capture, replicate, and scale expert-level reasoning across the workforce.

2.3 The Limits of Traditional Automation

Industrial automation has historically excelled at well-defined, deterministic tasks. Control loops maintain temperature and pressure. PLCs operate equipment sequences. MES systems enforce production logic. Alarm systems notify operators when conditions fall out of tolerance.

However, automation falls short in areas requiring:

- Interpretation of complex, multi-variable signals
- Context-aware decision-making
- Cross-system coordination
- Adaptation to changing plant conditions
- Multi-step workflows that involve humans and machines

For example, an automation system may detect a pressure spike, but it cannot correlate it with a slow drift in vibration, a clogged filter upstream, or a maintenance report from three months ago. Nor can it determine which technician is available, generate a work instruction tailored to the exact scenario, or update multiple systems with downstream actions.

Traditional automation is foundational—but insufficient. It is designed to control machinery, not interpret meaning or execute decisions across systems.

This is where agentic AI fills the gap.

2.4 Operational Complexity Is Increasing Faster Than Human Capacity

Modern industrial processes generate more data, operate under tighter tolerances, and require faster interventions than ever before. At the same time, the number of systems, interfaces, and workflows has multiplied.

Operators today must interpret:

- Thousands of real-time sensor feeds
- Complex causal relationships within processes
- Hundreds of alarms per shift
- Detailed maintenance histories
- Evolving production constraints
- Safety and compliance requirements
- Human factors, staffing, and shift changes

This overload can lead to delayed responses, incomplete analysis, or inconsistent decisions—especially during high-pressure events or abnormal operations.

Despite improvements in analytics tools, most organizations still rely heavily on manual decision-making. Humans must:

- Notice a deviation
- Investigate its root cause
- Consult documents, trends, and logs
- Decide what action to take
- Trigger work orders or workflows
- Communicate tasks to technicians
- Verify resolution

Each step introduces delays and variability. In this environment, even skilled teams struggle to maintain optimal performance.

Agentic AI is not about replacing humans—it is about reducing cognitive load, accelerating decision cycles, and ensuring actions are timely, consistent, and coordinated.

2.5 The Insight-to-Action Gap

Industrial companies have made significant progress in deploying dashboards, KPIs, predictive models, and advanced analytics. But the value of these systems depends on what happens after an insight is generated. Too often, insights are trapped in dashboards or reports, requiring humans to manually translate them into decisions.

This creates a critical bottleneck:

- Predictive analytics predicts a failure → but no action is taken
- Quality systems detect deviations → but root cause analysis takes hours
- Energy optimization models recommend adjustments → but operators lack bandwidth
- Alarms trigger dozens of alerts → but prioritization depends on individual judgment

The last mile—from insight to action—remains painfully manual.

Agentic AI can bridge this gap by turning insights into autonomous workflows, executed by intelligent agents that can coordinate across systems, people, and assets.

2.6 The Need for a New Operational Paradigm

The combined pressures of data complexity, workforce constraints, fragmented systems, reactive processes, and manual decision-making have pushed the industrial sector to an inflection point.

Incremental improvements are no longer enough. Organizations need a fundamentally new operational model:

- One where systems understand industrial context
- One where intelligence can take action autonomously
- One where workflows are executed consistently and efficiently
- One where knowledge is captured, scaled, and reused
- One where humans and AI collaborate seamlessly

This new model is made possible by industrial agents.

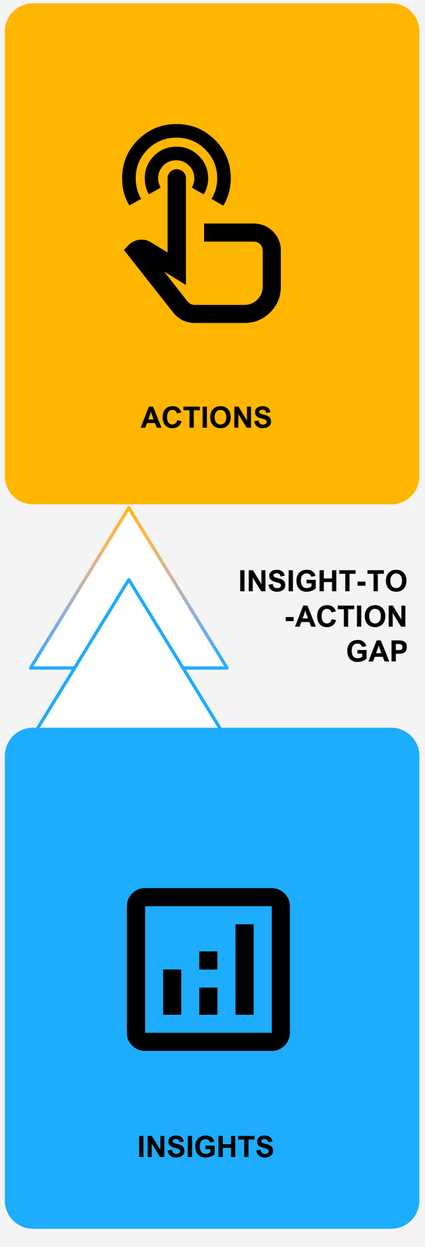

3. The rise of Agentic AI in industry

Industrial organizations have always depended on the coordination of people, machines, and information. Operators monitor equipment. Engineers troubleshoot anomalies. Maintenance teams interpret signals and logs. Supervisors orchestrate workflows across shifts. For decades, this human-centered model has been the backbone of industrial performance.

But today's industrial environment is too complex, too fast-moving, and too data-rich for human teams alone to manage optimally. Thousands of signals stream in every second. Hundreds of assets operate in interdependent sequences. Dozens of systems store slices of context. And operational risks—from downtime to safety incidents—can escalate within minutes.

These conditions demand an evolution in how intelligence is created, shared, and acted upon. The next leap forward is not simply more automation or better dashboards. It is a new class of intelligent systems that can understand industrial context, reason across data modalities, interact with human teams, and take timely action.

This paradigm is known as agentic AI.

3.1 What is Agentic AI?

Agentic AI refers to intelligent systems—called agents—that can:

- Observe their environment
- Interpret and reason about context
- Identify goals or constraints
- Make informed decisions
- Take autonomous actions
- Collaborate with other agents or humans
- Learn from outcomes to improve future behavior

In other words, agents move AI beyond passive pattern recognition and into active operational participation. They behave as digital teammates with assigned responsibilities, decision-making authority, and the ability to execute actions.

In an industrial context, this means agents can:

- Monitor asset health
- Detect anomalies
- Identify root causes
- Generate or execute workflows
- Cross-reference documents and P&IDs
- Guide frontline workers
- Update systems such as CMMS or MES
- Communicate with other agents to coordinate actions

This is a dramatic shift from traditional automation, which is rigid, rule-based, and siloed.

Agentic AI is dynamic, contextual, and collaborative.

3.2 Why Industry Is the Ideal Environment for Agentic AI

Industrial operations present challenges that align perfectly with the strengths of agentic systems.

High Data Volume and Variety

Plants generate billions of data points, from time-series to operational logs to engineering drawings. Agents thrive in data-rich environments.

Complex, Interconnected Processes

Issues rarely have a single cause. Agents can analyze multi-dimensional relationships using the industrial knowledge graph.

Critical Need for Speed and Precision

Anomalies can cascade quickly. Agents react instantly and consistently—far faster than human teams.

Structured Workflows and Procedures

Industrial processes already fit into logical structures ideal for agent execution.

High Cost of Error

Downtime, safety, compliance—agents reduce variability and prevent mistakes.

Thus, industrial operations are not only compatible with agentic AI—they *demand* it.

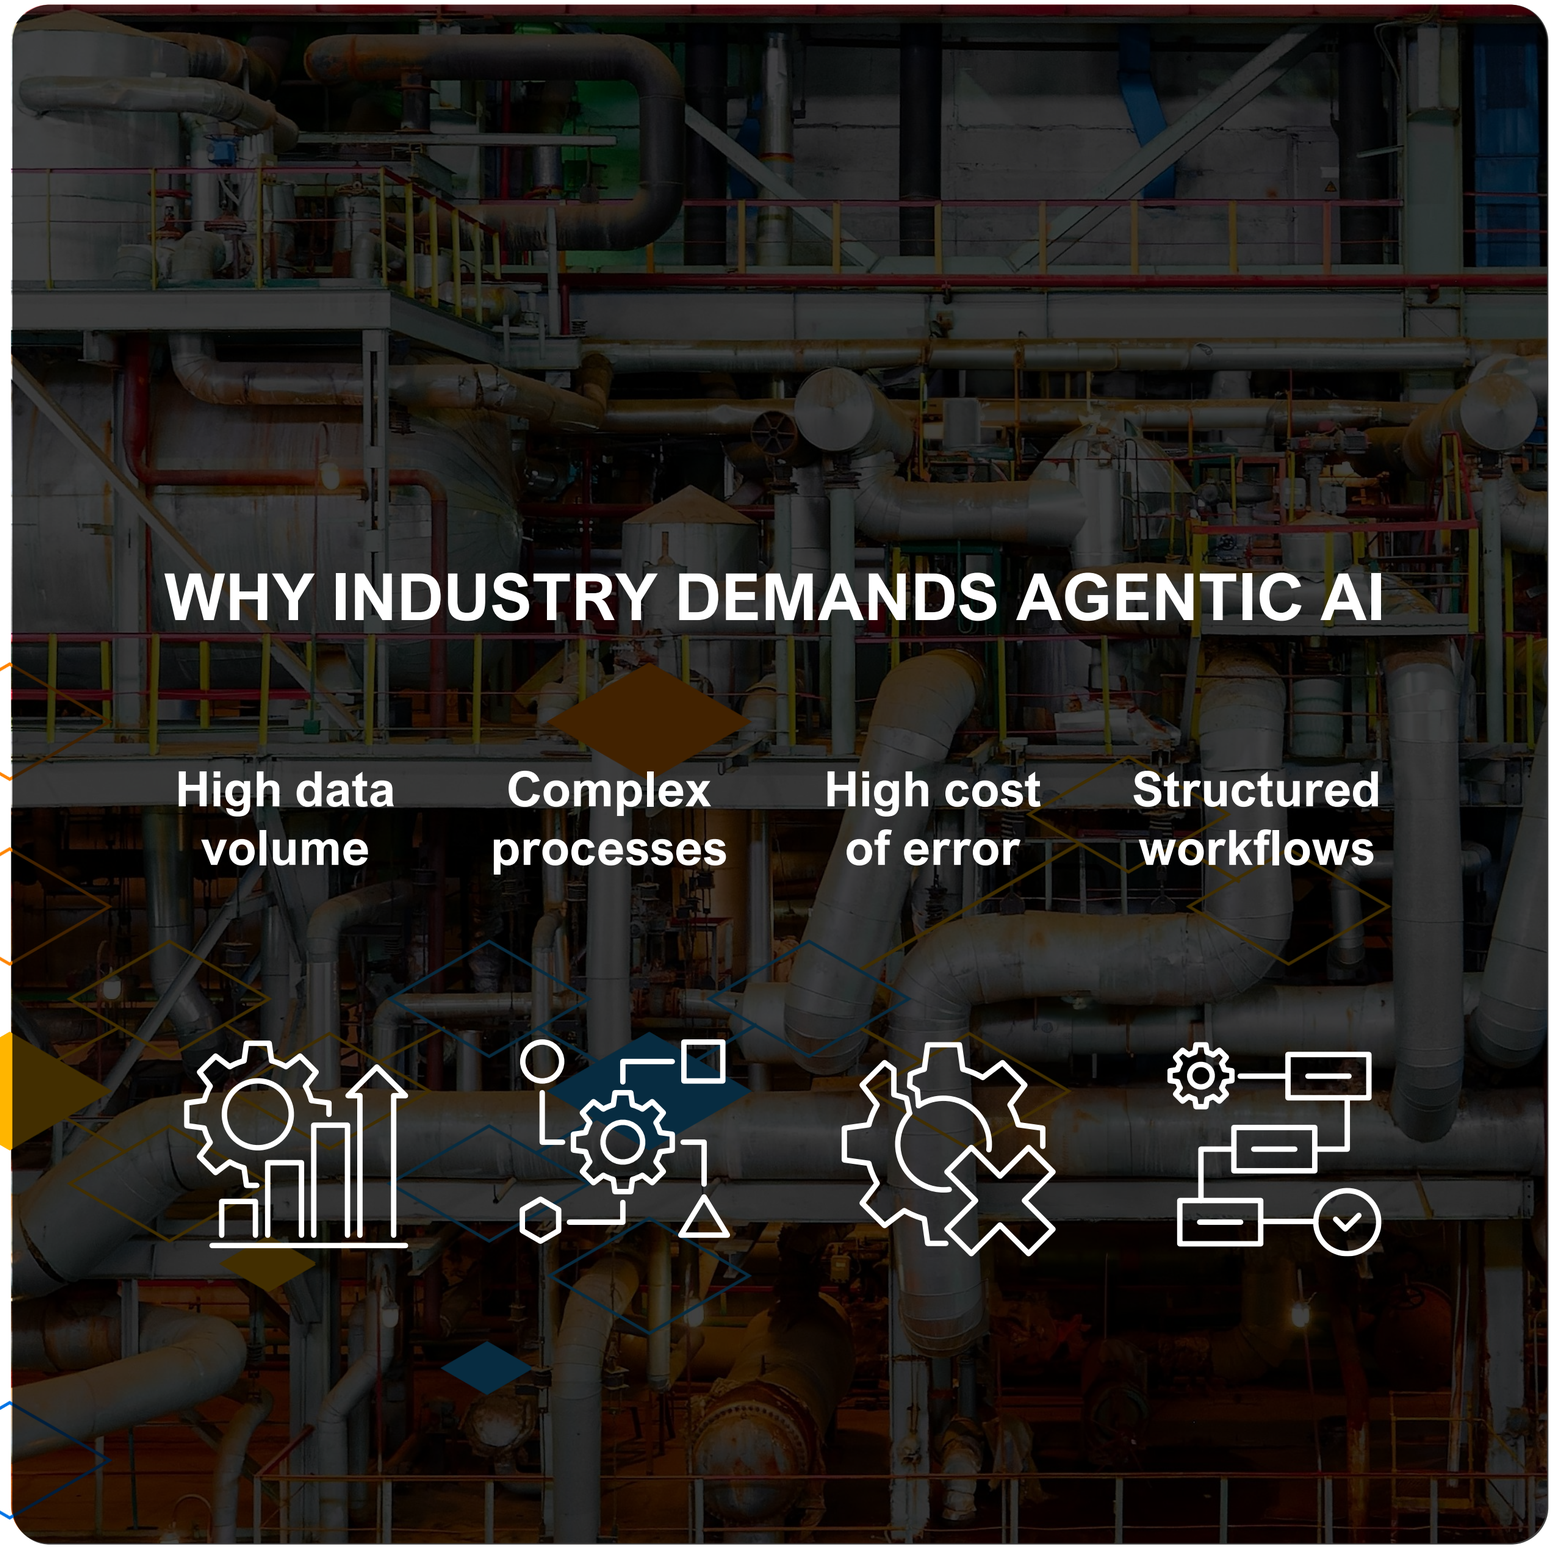

4. Foundations needed for industrial agents

Industrial agents represent a profound shift in how intelligence is applied across asset-intensive operations. But unlike consumer AI agents or generic multi-agent frameworks, industrial agents operate in environments where the consequences of error can be severe—ranging from equipment damage and unplanned downtime to environmental incidents and safety risks. This means industrial agents cannot be experimental toys or loosely governed chat-driven tools; they must be designed on a strong technical foundation that guarantees contextual accuracy, operational safety, transparency, and disciplined governance. In short, agentic AI in industry requires a deeper, more deliberate architectural underpinning than any other enterprise AI domain.

The foundational requirement for industrial agents is context. This cannot be overstated. Industrial operations generate a staggering variety of data—time-series signals, alarms, engineering drawings, procedural documentation, logs, work orders, event histories, shift notes, and more—but none of these data streams carry meaning on their own. A temperature reading, for example, has no useful interpretation unless the system understands which asset produced it, where that asset sits within the process, how it interacts with other equipment, what its normal ranges are, and how its behavior correlates with historical failures.

Without this contextual layer, even the most powerful AI models cannot reliably interpret the operational significance of a signal. That is why the first pillar for industrial agents is a robust semantic knowledge graph that unifies disparate data sources and encodes the relationships and semantics necessary for meaningful reasoning.

The semantic industrial knowledge graph provides this missing structure. It creates a living, interconnected representation of assets, subsystems, process flows, sensors, documents, and historical events. It does not simply store data—it represents meaning. Assets are connected to the signals they produce, the maintenance actions historically performed on them, the procedures that govern their safe operation, the flow networks they participate in, and the failure patterns that tend to manifest over time.

Engineers, technicians, and operations teams have always carried this context in their heads; the knowledge graph now provides it digitally, continuously, and at scale. For agentic AI, the knowledge graph serves as both the factual grounding and the navigation map for decision-making. It ensures agents respond with deterministic, reality-bound outputs rather than probabilistic guesses.

Layered on top of the knowledge graph is the Industrial LLM, which provides agents with the ability to interpret language, understand engineering concepts, and reason through technical scenarios. A generic LLM trained on internet data lacks the vocabulary and semiconductor-level intuition that industrial environments require. It cannot reliably differentiate between asset types, interpret a P&ID, or understand the implications of an upstream pressure drop on downstream quality. SymphonyAI's Industrial LLM, by contrast, is trained on engineering schematics, operating procedures, FMEAs, maintenance logs, equipment behavior patterns, and domain-specific documentation across manufacturing, energy, chemicals, metals, and other process industries.

This allows agents to understand and generate language consistent with industrial expertise and to draw from a reservoir of embedded knowledge that aligns with how plants and equipment function in reality.

Even with contextual structures and domain-specific intelligence, industrial agents still require a real-time, high-quality data backbone to operate effectively. Industrial environments are dynamic, and agents must be able to detect and respond to changes as they happen.

This necessitates an Industrial DataOps layer capable of ingesting millions of data points per second, contextualizing them in real time, enforcing data quality and lineage, and providing a unified namespace that bridges IT, OT, IoT, and engineering systems. Without clean, reliable, timely data, agents cannot differentiate between noise and meaningful deviations, nor can they take informed action. DataOps also ensures data is processed consistently across cloud and edge locations, enabling agents to operate seamlessly whether they are running inside a central system or on local edge nodes near the equipment they monitor.

Industrial agents must also operate within a rigorous governance framework that ensures safety, reliability, and accountability. Unlike traditional AI deployments, where outputs may simply be reviewed or ignored, industrial agents have the ability to trigger workflows, issue instructions, and initiate real-world actions. This requires a governance model that includes clear approval rules, audit trails, restricted action policies, safety interlocks, and role-based autonomy levels. Operators and engineers must be able to see why an agent recommended a particular action, validate the reasoning, and approve or decline interventions when necessary.

Governance also ensures organizations can control how autonomous an agent is allowed to become, ranging from assistive to augmented to fully autonomous roles. This structured autonomy is crucial to gaining trust and ensuring compliance with safety standards and regulatory requirements.

A further foundational element is multimodal AI capability—the ability to interpret not only text and numbers, but also documents, images, P&IDs, 3D models, and other industrial data formats. Industrial anomalies rarely present themselves in a single modality. For instance, a deviation in vibration might only be fully understood when correlated with shift logs, maintenance reports, engineering diagrams, and upstream flow data.

An agent must be able to navigate between these data types fluidly, extracting meaning from each and combining those insights into a coherent diagnosis or action plan. Without multimodal intelligence, agents operate with blind spots, and blind spots in industrial environments can lead to misinterpretation and suboptimal decisions.

Lastly, many industrial use cases require edge intelligence, where agents must operate close to the assets themselves.

High-speed control decisions, offline operation in remote locations, cybersecurity constraints, and low-latency requirements mean industrial agents cannot rely solely on cloud-based inference. By supporting edge deployment, agents can maintain operational continuity during network outages, respond in milliseconds to critical events, and integrate directly with OT systems such as PLCs, SCADA, and control loops. This edge-cloud hybrid model gives organizations the flexibility to distribute intelligence where it is needed most while maintaining centralized oversight and learning across all agents.

When combined, these foundations—semantic knowledge graphs, Industrial LLMs, real-time DataOps, rigorous governance, multimodal AI, and edge intelligence—form the essential stack for agentic AI in industrial environments. Without these capabilities, agent systems cannot be trusted to operate safely or intelligently; with them, they become powerful collaborators capable of transforming operations from reactive to autonomous.

These requirements also highlight why generic AI tools are insufficient for high-stakes industrial use. Industrial agents cannot be trained on internet data, reason without grounding, or make decisions based on approximations

They must understand how a plant works, how assets interact, how safety is ensured, and how workflows unfold across shifts and systems. They must be context-aware, failure-aware, procedure-aware, and outcome-aware. This is a level of precision only possible when all foundational components are integrated into a cohesive platform.

.

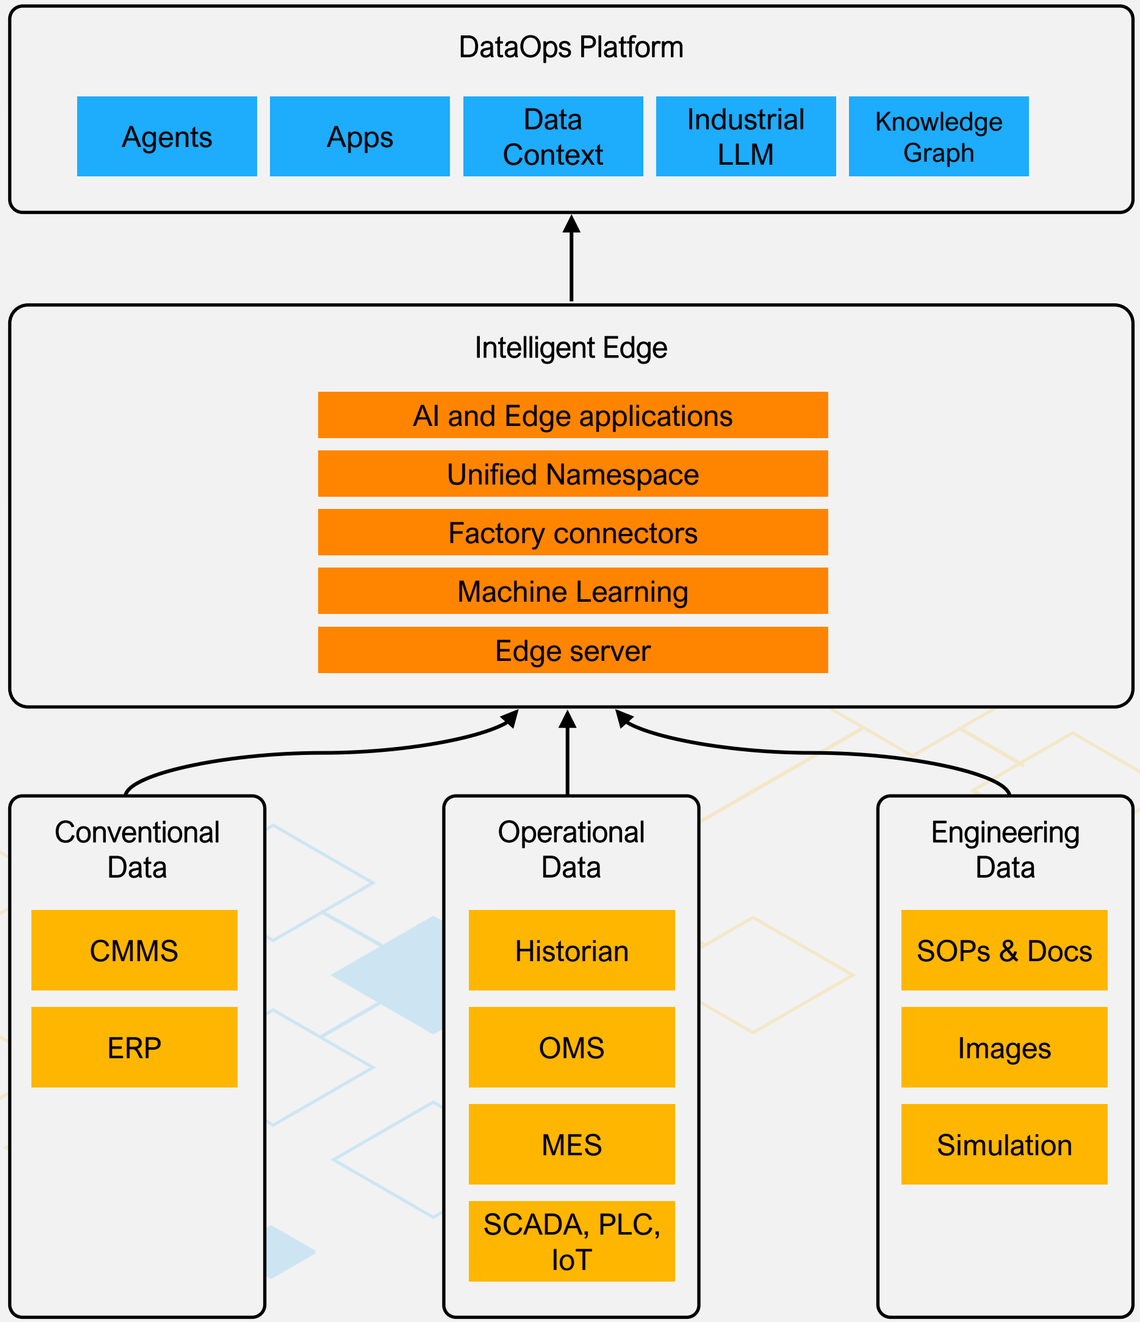

5. IRIS Foundry: the platform that makes Agentic AI possible

Industrial agents cannot operate on generic data platforms or loosely connected AI tools. They require a foundation engineered for the realities of industrial complexity—a platform capable of delivering context, intelligence, governance, multimodal understanding, and seamless integration across OT, IT, engineering, and environments. IRIS Foundry is this platform. It is the engine that powers IRIS Flows, enabling industrial agents to reason safely, act autonomously, and collaborate intelligently. IRIS Foundry unifies data, intelligence, and action into a single operational fabric. The following sections break down the foundational components that make IRIS Foundry uniquely suited to support agentic AI at industrial scale.

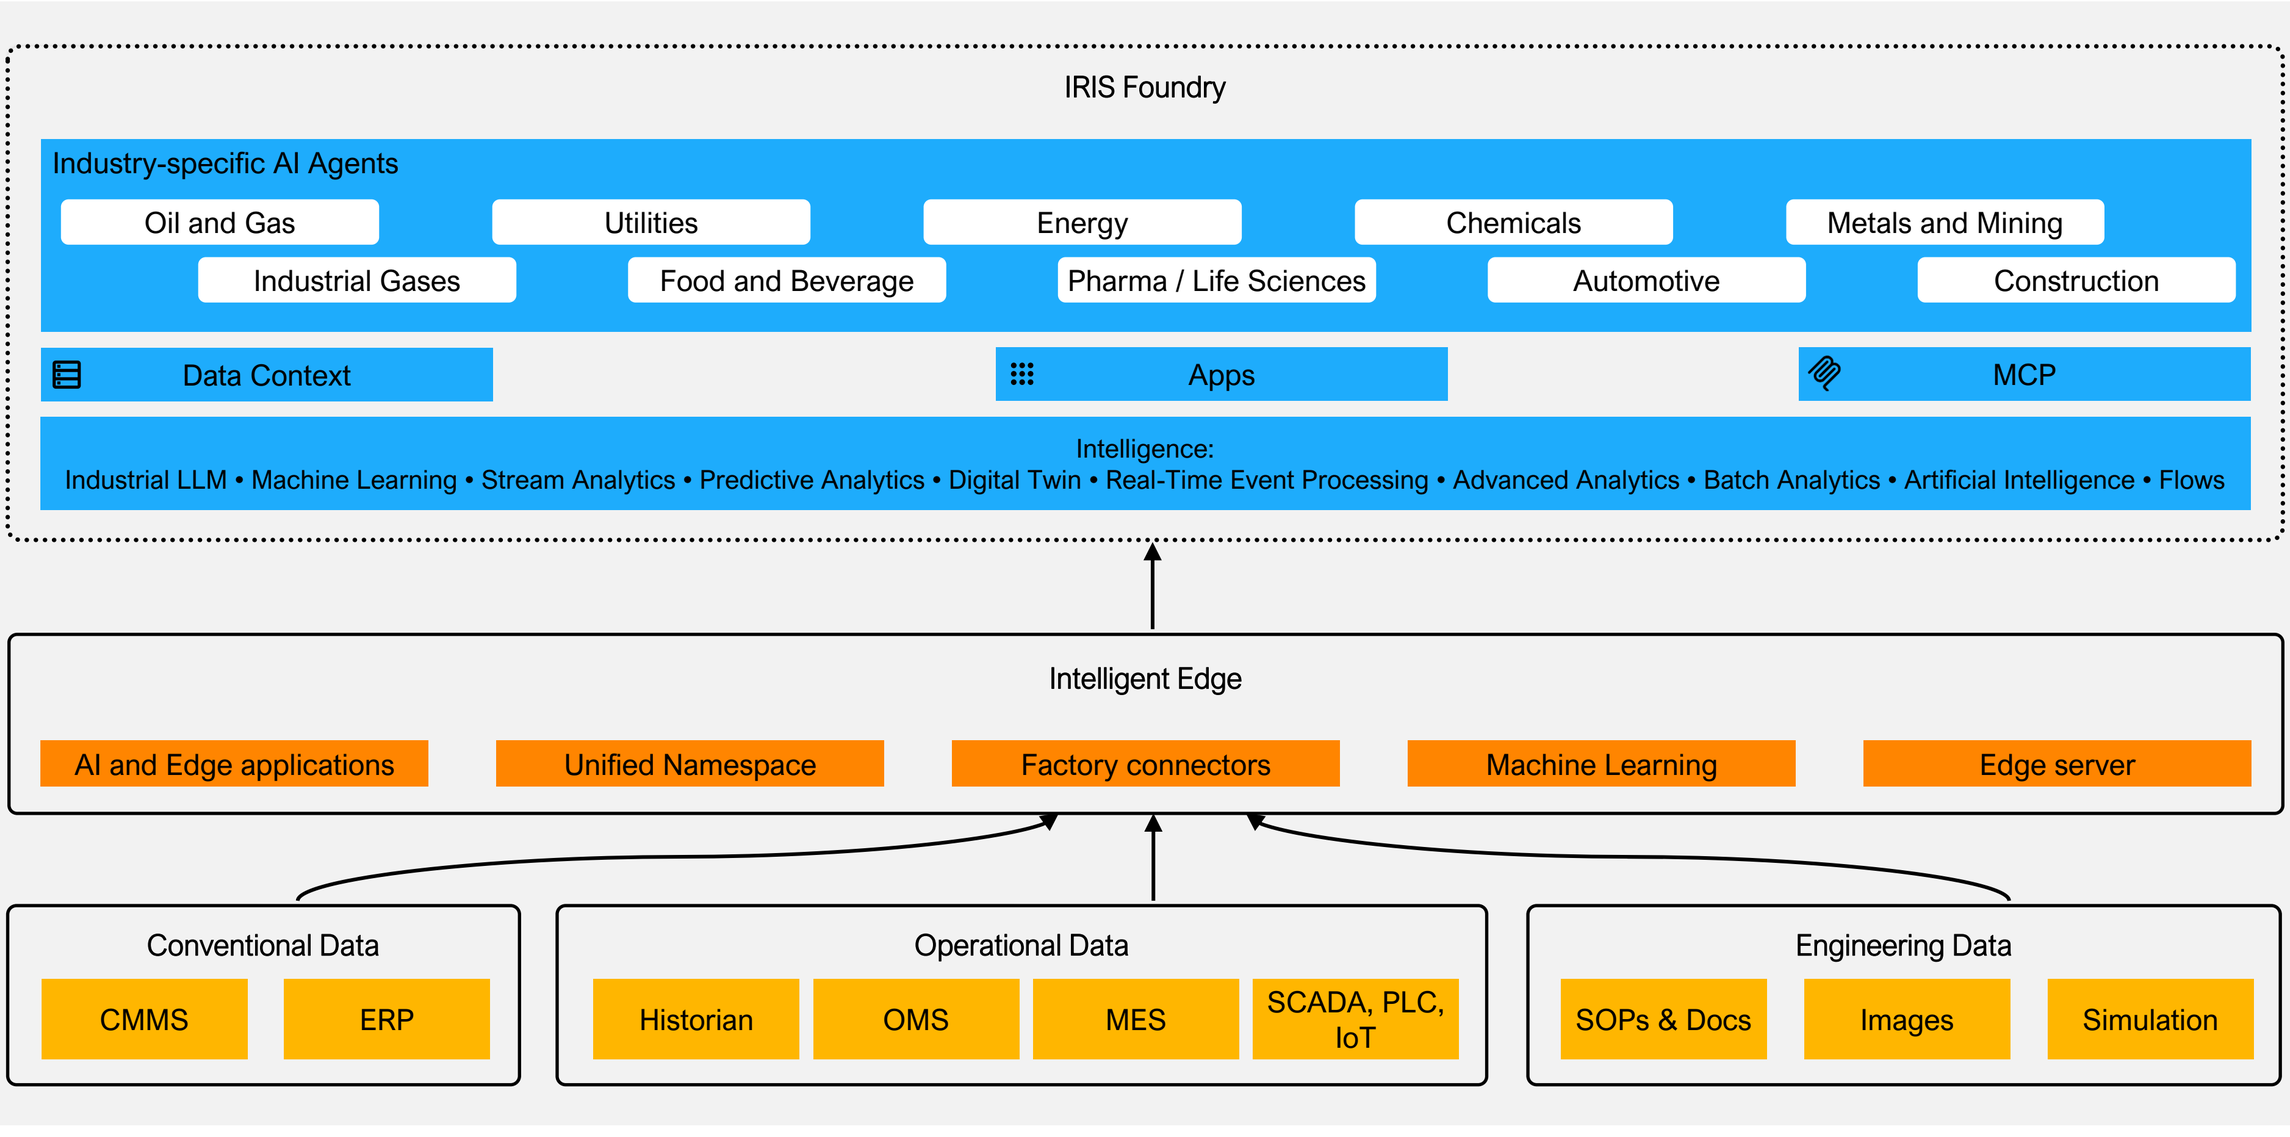

5.1 The Semantic Industrial Knowledge Graph: The Context Engine

At the core of IRIS Foundry lies its semantic industrial knowledge graph—a dynamic, interconnected representation of everything that defines plant operations. It captures the relationships among assets, processes, sensors, documentation, maintenance histories, operational events, and engineering models. Unlike a database, the knowledge graph encodes meaning, allowing agents to understand how different elements relate within the physical world.

For example, when an agent detects a pressure spike, the knowledge graph enables it to trace upstream and downstream dependencies, correlate historical failure patterns, interpret associated work orders, identify which SOP governs the situation, and determine how the event impacts production. This contextual fabric eliminates guesswork and ensures agents operate with the same holistic awareness that experienced operators develop over decades.

The knowledge graph is continuously updated, meaning agents always have access to the most current operational context. It serves as the navigational map for all agent reasoning and is what transforms data into actionable intelligence.

5.2 The Industrial LLM: The Domain Intelligence Layer

IRIS Foundry is powered by SymphonyAI's Industrial LLM, a large language model trained not on generic internet text but on the vocabulary, semantics, and logic that govern industrial operations. This includes engineering schematics, equipment manuals, P&IDs, FMEAs, procedures, maintenance logs, shift handover notes, and process documentation from asset-intensive industries.

This domain-rich training gives the Industrial LLM a profound understanding of how plants operate, how assets behave, how failures manifest, and how humans articulate operational tasks. Agents built on top of this model can interpret complex questions, summarize dense engineering documents, explain anomalies, and generate accurate troubleshooting guidance.

More importantly, the Industrial LLM is tightly coupled with the knowledge graph, ensuring its reasoning is grounded in deterministic truth rather than probabilistic hallucination. This synergy gives IRIS agents the ability to respond intelligently to domain-specific queries—and to do so safely.

5.3 Industrial DataOps: The Unified, Real-Time Operational Backbone

A knowledge graph and Industrial LLM are only as useful as the data they understand. This is why IRIS Foundry includes a fully integrated Industrial DataOps architecture designed for the ingest, transformation, contextualization, and governance of massive volumes of industrial data.
IRIS Foundry can handle millions of data points per second from OT systems, IoT devices, engineering repositories, historians, SCADA, MES, CMMS, ERP, and more. Rather than simply aggregating data, IRIS Foundry contextualizes it in real time—linking tags to assets, linking events to equipment, aligning telemetry with operational conditions, and validating data quality continuously.

This ensures that agents in IRIS Flows always base their reasoning on clean, trusted, real-time data, eliminating the risks associated with stale or inconsistent information. The DataOps layer also provides lineage, auditing, schema harmonization, and automated data quality checks to ensure compliance and reliability across the enterprise.

5.4 Edge-to-Cloud Intelligence: Consistent Execution Everywhere

Industrial environments often span multiple sites, each with varying connectivity, latency requirements, and system constraints. IRIS Foundry supports a hybrid architecture where intelligence can run seamlessly at both the edge and in the cloud.
At the edge, IRIS Foundry enables near-real-time decision-making for tasks such as anomaly detection, local asset monitoring, and immediate workflow triggering—even in offline or bandwidth-constrained contexts.

In the cloud, IRIS Foundry orchestrates multi-agent collaboration, enterprise-wide optimization, cross-site learning, and global intelligence models.
This hybrid architecture ensures agents remain operational, responsive, and context-aware regardless of infrastructure constraints. It is essential for enabling IRIS Flows to function in the real world, where milliseconds matter and network interruptions cannot be allowed to halt critical decision-making.

5.5 Governance, Security, and Trust: The Guardrails for Autonomy

Industrial autonomy is only valuable if it is safe. IRIS Foundry includes robust governance, security, and compliance frameworks that control how agents behave, which actions they are allowed to take, and under what conditions they require human approval.

Every data point, decision, and action is tracked through audit trails. Agents operate under defined autonomy levels—assistive, augmented, or autonomous—with role-based controls determining who can initiate, approve, or override agent actions. Comprehensive identity management, encryption, access permissions, and compliance policies ensure agents act within strict operational boundaries.

This governance structure builds trust—operators, engineers, and leadership must be confident that agents will behave predictably, transparently, and safely.

IRIS Foundry provides that confidence through rigorous safeguards, explainability mechanisms, and continuous monitoring of agent behavior.

5.6 Multimodal AI Capabilities: Understanding Every Dimension of Operations

Industrial intelligence cannot rely on a single mode of data. A plant's reality lives simultaneously in sensor streams, engineering documents, maintenance histories, 3D models, shift notes, and human inputs. IRIS Foundry supports multimodal AI that can interpret all of these formats cohesively.

Agents can read P&IDs, analyze time-series signals, correlate deviations with historical logs, extract procedures from manuals, and visually interpret 3D representations of assets.

This multimodal grounding enables agents to diagnose issues, propose interventions, and execute workflows with a level of sophistication that mirrors expert human reasoning.

This ability to draw from multiple data modalities is what allows agents to understand complex industrial scenarios, where the root cause of an event may span several systems and data types.

5.7 The SymphonyAI Ecosystem: Closing the Loop from Insight to Action

IRIS Foundry integrates seamlessly with the broader SymphonyAI Industrial ecosystem, enabling agents not only to analyze and decide but also to act. IRIS Copilot provides a natural language interface where operators can query agents, validate reasoning, or request assistance.

Proceedix Connected Worker brings intelligence directly to the field, enabling agents to deliver AR-guided procedures, inspection workflows, and safety instructions to technicians. IRIS Forge empowers teams to build no-code applications that incorporate agent intelligence into dashboards, mobile experiences, and shift operations.

This ecosystem allows IRIS Flows agents to close the loop entirely—from insight to action to verification—bridging the last mile that often breaks in traditional industrial AI deployments.

5.8 Enterprise Scalability: Built for Multi-Site, Multi-Agent Operations

Finally, IRIS Foundry provides enterprise-grade scalability that supports thousands of agents across global operations. Whether deployed in a single plant or across a network of sites worldwide, IRIS Foundry maintains consistency, synchronization, and performance. It manages massive data volumes, orchestrates distributed intelligence workloads, and ensures agents always have access to the full operational context they need to make informed decisions.

This scalability is essential for organizations seeking to build a digital industrial workforce—a coordinated, intelligent layer of agentic AI that complements and enhances human teams across the entire enterprise.

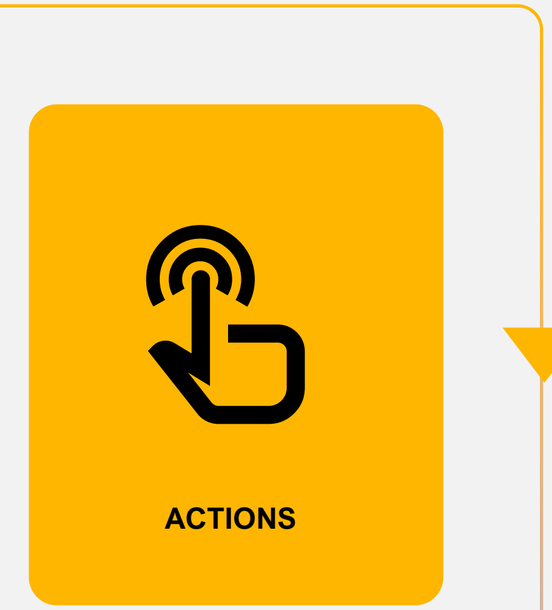

The SymphonyAI Ecosystem:
Closing the Loop from Insight to Action

6. IRIS Flows: a frontier in industrial Agentic AI

The emergence of industrial agents marks a pivotal moment in the evolution of operational intelligence. For decades, industrial facilities have relied on supervisory control systems, predictive analytics, dashboards, and human expertise to interpret data and drive decisions. But as plants have grown more complex, more automated, and more data-dense, the gap between raw information and effective action has widened. Operators are overloaded. Systems remain siloed. Insights often fail to translate into timely execution. And much of the knowledge that keeps operations stable lives inside the minds of a shrinking pool of experienced personnel.

IRIS Flows is SymphonyAI's answer to this challenge—a transformative environment designed to bring industrial agents to life. It provides a unified place where intelligent agents can be created, deployed, orchestrated, and governed, all while operating on the rich contextual understanding of IRIS Foundry. IRIS Flows is not another automation engine or workflow tool; it is a fully agentic execution layer that enables autonomous, goal-driven behavior across processes, equipment, and teams.

Through IRIS Flows, organizations can build a digital industrial workforce—a coordinated system of agents that continuously monitor conditions, diagnose issues, guide frontline workers, trigger automated workflows, and learn from outcomes.

6.1 The Philosophy Behind IRIS Flows: From Tasks to Teams

Traditional automation focuses on tasks: If X happens, do Y. The complexity of industrial operations makes this insufficient, because real-world events rarely follow linear patterns. Root causes evolve, conditions shift, and outcomes depend on multi-variable dependencies.

IRIS Flows is built around a different philosophy—teams instead of tasks. Each agent is assigned a responsibility rather than a fixed instruction. One agent may track the health of a rotating asset. Another may evaluate energy consumption patterns. A third may interpret work order data. A fourth may guide technicians through a repair procedure. Together, these agents behave like a digitally augmented team, collaborating to maintain optimal performance across operations.

This concept of "digital teamwork" is foundational. Agents do not operate in silos.

These agents do not function as isolated scripts or bots, but as interconnected collaborators with shared context, shared awareness, and a unified operational purpose.

They communicate, delegate, escalate, and cooperate, mirroring real-world human coordination but with the speed and consistency of AI.

6.2 The IRIS Flows Studio: Creating Industrial Agents Without Code

IRIS Flows includes a visual, intuitive environment where industrial teams can build agents using natural language and drag-and-drop components. This lowers the barrier to entry and makes industrial agent development accessible to reliability engineers, operations leaders, maintenance planners, and process specialists—people who understand the plant deeply but do not write code.

In IRIS Flows Studio, users can:

- Describe the agent's goal in natural language
- Define the conditions it should monitor
- Connect it to relevant real-time and historical data
- Specify the types of actions it is allowed to take
- Configure safety and approval boundaries
- Orchestrate interactions with other agents

The environment automatically translates this intent into executable logic powered by the Industrial LLM, the knowledge graph, and IRIS Foundry's DataOps underpinnings.

This natural-language-first creation model makes agent development fast, safe, and domain-aligned.

Just as importantly, IRIS Flows Studio supports iteration. Teams can observe how agents behave, refine their responsibilities, adjust decision parameters, and expand their capabilities gradually over time. This evolutionary approach mirrors how human expertise develops and ensures agents become more effective as their environment changes.

6.3 A Library of Prebuilt Industrial Agents

To accelerate adoption, IRIS Flows includes an extensive library of more than 300 prebuilt industrial agents covering a wide range of operational domains.

These prebuilt agents offer immediate value by encapsulating validated logic for:

- Reliability and predictive maintenance
- Anomaly detection and multivariate signal interpretation
- Alarm management and prioritization
- Root cause analysis
- Quality and deviation analysis
- Energy and emissions optimization
- Work order generation and enrichment
- Safety and compliance monitoring
- Shift operations and reporting
- Document intelligence and SOP retrieval
- Procedure generation and connected worker routing

Each prebuilt agent is fully customizable, allowing organizations to rapidly deploy out-of-the-box intelligence while tailoring behavior to site-specific needs. These agents were designed based on industry best practices and refined through real-world deployments, providing a fast path to value while building confidence in the agentic paradigm.

6.4 Connected Agents: How IRIS Flows Enables Multi-Agent Collaboration

Industrial operations are inherently interdependent. No single agent can understand, optimize, or control an entire system. Instead, agents must collaborate—sharing insights, delegating tasks, and synchronizing workflows. IRIS Flows orchestrates this collaboration efficiently and safely.

For example, consider an event such as a temperature increase on a critical piece of rotating equipment:

- A monitoring agent detects the anomaly.
- A reasoning agent correlates temperature changes with changes in vibration patterns and throughput.
- A diagnostic agent identifies the most likely cause based on historical failure modes and documentation.
- A maintenance agent evaluates whether the issue warrants intervention and, if so, initiates a work order.
- A connected worker agent prepares a context-specific inspection procedure, pulling from manuals and prior incidents.
- A verification agent confirms post-maintenance metrics have returned to normal.
- A learning agent logs the incident and updates the knowledge base for future predictions.

These agents behave like members of a well-coordinated operational team—each contributing its expertise, each operating with full awareness of the others, and each aligned to plant-level goals.

IRIS Flows handles the communication, sequencing, and decision boundaries that allow this collaboration to occur seamlessly.

6.5 Agent Autonomy Models: Assistive, Augmented, and Autonomous

Not every agent should act autonomously, and not every organization is ready for high levels of agent autonomy on day one. IRIS Flows supports multiple autonomy models to match organizational readiness, regulatory requirements, and safety constraints.

In assistive mode, agents function as copilots: they monitor, analyze, and recommend, but require human approval before taking any action. This mode is ideal for early deployments and high-risk decisions.

In augmented mode, agents execute actions automatically but only in non-critical operations or with predefined boundaries. Humans intervene only when an agent flags ambiguity or risk.

In fully autonomous mode, agents operate independently within well-defined constraints. They are empowered to diagnose, act, update systems, and coordinate workflows without requiring human input—ideal for routine or low-risk processes.

These graduated autonomy levels help organizations build trust, validate performance, and expand capabilities in a controlled manner.

6.6 Human and Machine Collaboration Through Connected Worker Integration

IRIS Flows doesn't replace workers—it empowers them. Through seamless integration with Proceedix Connected Worker, agents can deliver step-by-step instructions, augmented reality overlays, and SOP-driven guidance directly to frontline technicians.

Operators receive:

- Dynamic procedures tailored to the asset's current condition
- AR visualizations showing exactly where and how to intervene
- Real-time answers to questions asked in natural language
- Instant access to relevant documents, diagrams, and manuals
- Intelligent escalation paths when discrepancies arise

This creates a true closed loop from digital detection to physical execution, with humans and agents working in orchestrated harmony.

6.7 Closing the Loop: How IRIS Flows Executes End-to-End Workflows

Traditional AI solutions often break down at the last mile—insight without action. IRIS Flows eliminates this bottleneck by enabling agents to execute full operational workflows spanning OT, IT, and frontline teams.

Agents can:

- Trigger work orders in CMMS systems
- Update asset health scores
- Modify MES or historian records based on validated observations
- Send instructions through Teams or mobile devices
- Initiate inspections through Proceedix
- Alert supervisors when safety risks are detected
- Generate reports and log learnings automatically

This closed-loop capability ensures operational improvements are not left to chance or manual interpretation. Every step—from anomaly detection to remediation—is handled with consistency, speed, and intelligence.

6.8 Safety, Governance, and Operational Trust

IRIS Flows is designed for use in critical environments, so safety and governance are fundamental. Every agent action is governed by explicit rules, guardrails, approval workflows, and auditability. Agents must justify their recommendations, operators must understand the rationale, and all actions are traceable.

Organizations retain total control over what agents can do, including:

- Which systems they can modify
- What types of actions require approval
- What thresholds trigger escalation
- Who is notified when interventions occur
- How autonomous each agent is allowed to be

This governance ensures agentic AI enhances operations without introducing risk.

6.9 IRIS Flows as the Digital Workforce Operating System

Taken together, IRIS Flows serves as the operational layer where industrial agents live, think, collaborate, and act. It translates data into insight, insight into workflows, and workflows into real-world outcomes.

It enables organizations to create a scalable digital workforce—not one agent replacing one human task, but a coordinated network of agents augmenting entire teams and processes.

IRIS Flows is the environment where industrial expertise becomes digital, where decisions are executed with consistency and intelligence, and where the future of autonomous operations begins.

IRIS Flows as the Digital Workforce Operating System

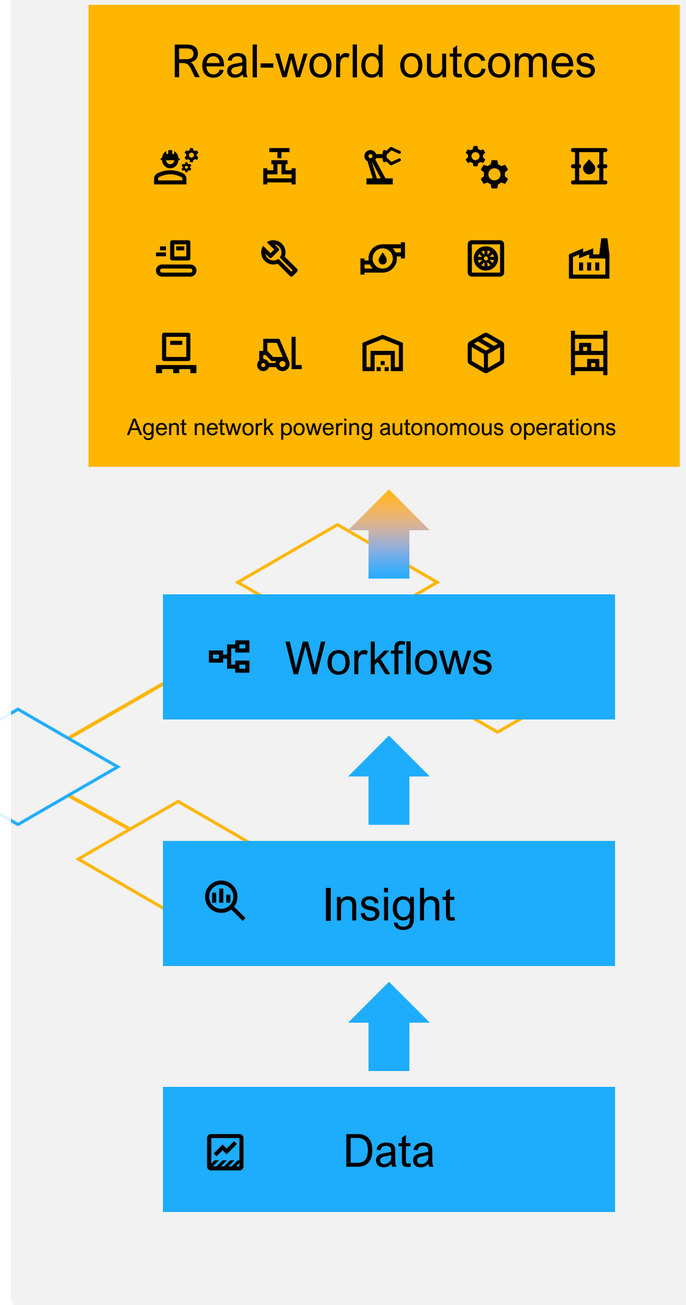

7. How IRIS Flows works

IRIS Flows represents a shift from applications that assist industrial operations to a system that can operate with intelligent autonomy. To achieve this, the architecture behind IRIS Flows must support a rich set of capabilities—real-time data ingestion, domain reasoning, agent orchestration, workflow execution, multimodal intelligence, and enterprise governance. Unlike conventional automation or AI platforms, IRIS Flows is designed to interpret context, coordinate multiple intelligent processes simultaneously, and act reliably within critical industrial environments.

At a high level, IRIS Flows consists of three foundational execution layers:

- The Agent Reasoning Layer,
- The Agent Orchestration Layer, and
- The Workflow & Action Execution Layer.

Surrounding these layers are key supporting capabilities including governance, data access, shared memory, safety boundaries, and integration services. Together, these components form an architecture built for real-world operational autonomy.

7.1 The Agent Reasoning Layer: Intelligence Powered by IRIS Foundry

At the heart of IRIS Flows lies the Agent Reasoning Layer, where each agent interprets data, evaluates context, and forms decisions. This layer is powered directly by the Industrial LLM, the semantic knowledge graph, and IRIS Foundry's multimodal data access.

When an agent is triggered—whether by a real-time signal deviation, a scheduled evaluation, or a request from another agent—it enters a reasoning cycle. During this cycle, it:

- Retrieves context from the industrial knowledge graph
- Accesses real-time and historical data from IRIS Foundry
- Interprets documents, diagrams, or logs related to the problem
- Applies domain-specific reasoning through the Industrial LLM
- Evaluates constraints such as safety limits or operational envelopes
- Generates an outcome—a recommendation, a diagnosis, a workflow step, or an action

The reasoning layer is intentionally modular. Each agent is composed of skills—discrete, reusable reasoning primitives—that allow for flexible and adaptive decision-making. One skill may focus on anomaly detection; another may specialize in extracting procedures from manuals; a third may evaluate KPI impacts. Skills can be combined, extended, or replaced, enabling organizations to evolve agent capabilities over time.

7.2 The Agent Orchestration Layer: Coordinating Teams of Agents

While individual agents can perform meaningful tasks, true operational transformation emerges when they collaborate. The Agent Orchestration Layer is the architectural component that choreographs how agents interact.

This layer handles:

- Inter-agent communication — where agents exchange insights or ask each other for assistance
- Delegation — where one agent assigns part of its responsibility to another
- Conflict resolution — ensuring two agents do not take contradictory actions

- Prioritization — determining which agent should take the lead in a multi-agent scenario
- Shared context memory — storing the evolving state of a workflow or event
- Supervisor agents — higher-level orchestrators that ensure alignment to operational goals

This orchestration layer turns a collection of intelligent processes into a coordinated system of digital teammates.

Consider a scenario where a pump shows abnormal vibration. One agent may detect the anomaly, but others may:

- Evaluate its severity,
- Diagnose the root cause,
- Check for prior related work orders,
- Assess the impact on production,
- Generate an inspection plan, and
- Validate the outcome.

The orchestration layer synchronizes these actions so they occur at the right time, in the right order, with the right safety boundaries.

Coordinated agent behavior is one of IRIS Flows' defining capabilities—and one of the primary reasons it can tackle complex industrial scenarios with high precision.

7.3 The Workflow & Action Execution Layer: Turning Decisions into Real Outcomes

Once agents reach a decision, IRIS Flows must translate that decision into action. The Workflow & Action Execution Layer is responsible for carrying out the tasks defined by agents and ensures that their recommendations manifest in real-world operational outcomes.

This execution layer includes integrations with:

- CMMS and EAM systems (work order creation and updates)
- MES and historian systems (record updates and operational adjustments)
- Alarm management platforms
- ERP and production planning systems
- Ticketing and collaboration tools such as Teams
- Proceedix Connected Worker for guided frontline execution

Actions may include initiating workflows, launching inspections, escalating issues to supervisors, sending alerts, updating system records, or executing model-driven optimizations.

Importantly, every action is subject to approval gates, safety constraints, and defined permissions.

IRIS Flows ensures that no agent can exceed its operational boundaries, and every step is logged for traceability and compliance.

7.4 Shared Memory and Context Persistence

Industrial events often unfold over hours, days, or even weeks. IRIS Flows includes a shared memory substrate that allows agents to track the state of ongoing events, workflows, and interventions.

This shared memory ensures that:

- Agents entering a scenario mid-process have full awareness of prior steps
- Multi-shift operations remain coherent across handovers
- Learnings from one event can influence future behavior
- Long-running workflows remain contextually intact

It provides continuity—something human workflows rely on heavily, but digital systems have historically lacked.

7.5 Real-Time Triggers, Event Processing, and Data Subscriptions

IRIS Flows supports true reactive intelligence through its real-time trigger and event-processing architecture, enabling agents to activate instantly in response to changing operational conditions. Whether a sensor crosses a critical threshold, another agent detects an emerging pattern, an alarm storm unfolds, a maintenance activity is logged, a quality deviation appears, a scheduled evaluation occurs, an operator submits a query, or an external system updates its state, each event carries contextual information directly into the agent reasoning layer.

This seamless flow of real-time signals allows agents to react within milliseconds, ensuring fast, context-aware decisions. By combining continuous monitoring with on-demand responsiveness, IRIS Flows adapts effortlessly to both steady-state operations and sudden, time-sensitive events, delivering intelligence that is always aligned with the pace and demands of industrial environments.

7.6 Edge Execution: Bringing Intelligence to the Plant Floor

Certain industrial decisions—such as identifying a rapidly developing anomaly or responding to equipment behavior that changes within milliseconds—must occur directly at the source. IRIS Flows enables agents to run seamlessly both in the cloud and at the edge, ensuring intelligence is available wherever it is needed most.

At the edge, agents can monitor equipment locally, perform safety-critical analyses, trigger alarms or instructions, and continue operating even when network connectivity is limited or unavailable. This low-latency, near-machine intelligence ensures real-time responsiveness, while cloud-based agents handle more computationally intensive tasks such as collaborative reasoning, multi-agent orchestration, and long-horizon optimization.

Together, this hybrid model provides the best of both worlds—ultra-fast decision cycles at the plant floor and enterprise-wide intelligence and coordination in the cloud.

7.7 Safety, Governance, and Agent Boundaries

Autonomy in industrial environments is only valuable when it operates within strict and transparent boundaries. IRIS Flows incorporates a comprehensive governance framework that dictates how agents behave, how much independence they have, and how every decision is validated. This includes clearly defined autonomy levels, structured approval workflows, explicit allowlists and denylists for agent actions, and safety envelopes that ensure decisions remain within acceptable operating limits. Every action taken by an agent is logged and auditable, and policies are enforced to ensure regulatory compliance and operational discipline.

Human override capabilities remain central, giving operators full control whenever necessary. These safeguards ensure agents not only make intelligent decisions, but also act predictably, safely, and in alignment with both plant-level procedures and enterprise governance standards.

7.8 High Availability and Multi-Site Scalability

IRIS Flows is engineered from the ground up to support the scale, resilience, and geographic diversity of global industrial enterprises. Its architecture incorporates mechanisms for horizontal scaling, load balancing, and distributed deployment across multiple plants, control systems, regional edge clusters, and redundant cloud environments. This ensures that large fleets of agents—often operating in the thousands—can run simultaneously without degradation, share intelligence across facilities, and maintain coherence even across

complex multi-region networks. Whether deployed in a single refinery or across dozens of manufacturing plants worldwide, IRIS Flows provides a robust, high-availability foundation that ensures agentic intelligence is always accessible, always reliable, and always aligned across the entire enterprise.

Examples of Industrial Agents

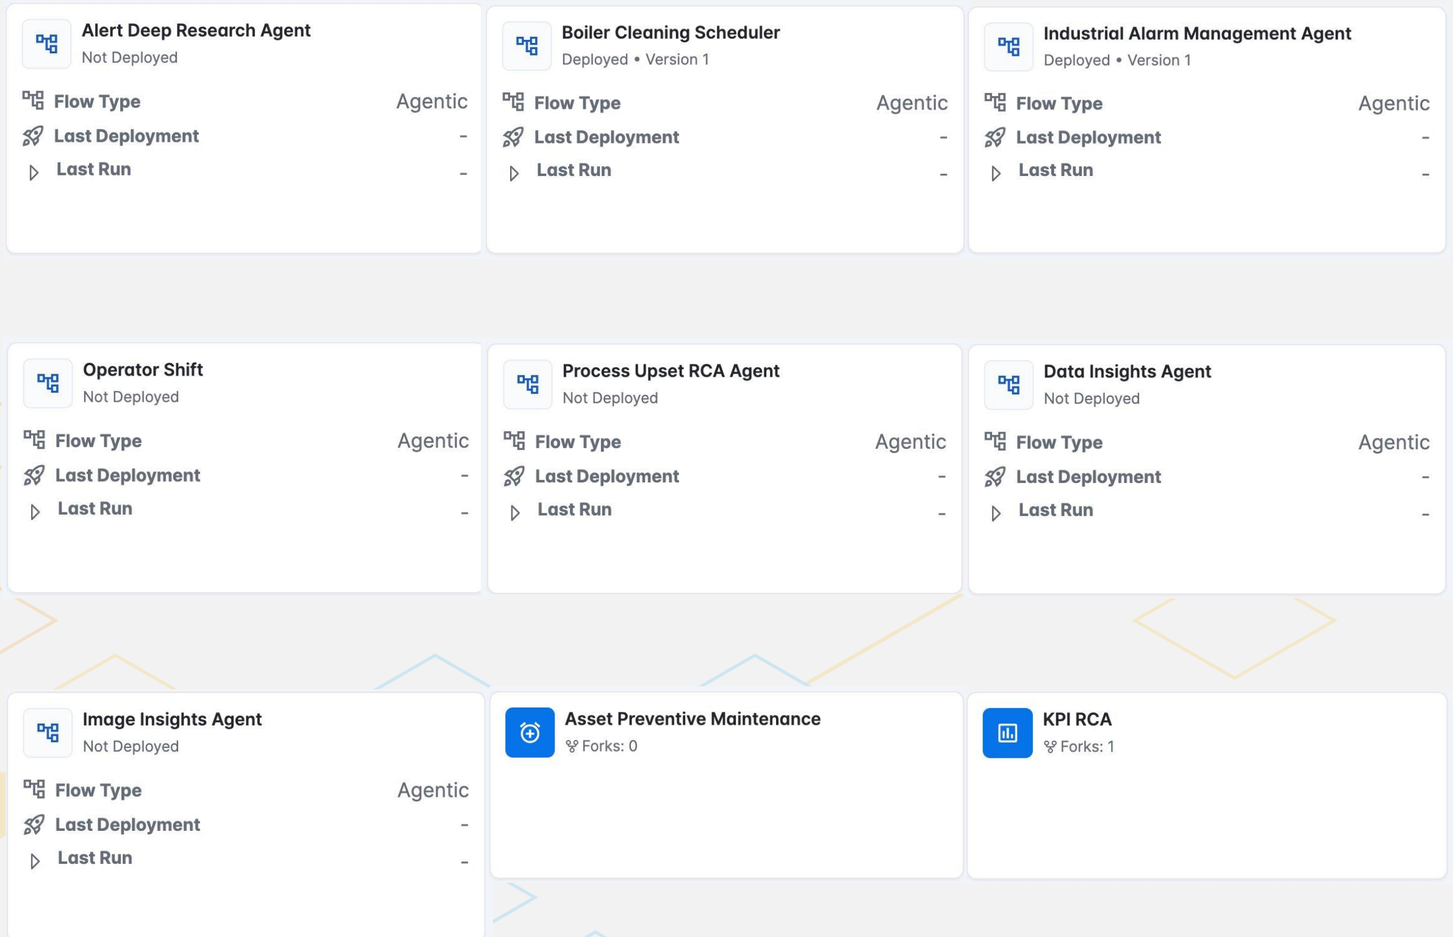

8. The IRIS Flows agent ecosystem

Industrial operations are too complex for any single agent or system to manage effectively. Problems rarely appear in isolation, and neither do their solutions. A vibration anomaly may relate to upstream flow restrictions, recent maintenance actions, environmental factors, or operator interventions. A quality deviation might stem from a combination of process drift, equipment wear, ambient conditions, or inconsistencies in procedures. Addressing such problems requires multiple perspectives, multiple skills, and coordinated action—much like a well-functioning human team.

IRIS Flows embraces this reality by enabling a diverse ecosystem of industrial agents, each with specialized abilities, responsibilities, and domain understanding. These agents collaborate seamlessly, sharing context, exchanging insights, and coordinating tasks to maintain optimal operational performance. This multi-agent ecosystem is a cornerstone of agentic AI and a defining capability of IRIS Flows.

8.1 The Role of Specialized Agents in Industrial Operations

Specialization drives efficiency in human teams, and the same principle applies to digital workforces. Within IRIS Flows, agents are intentionally designed with focused roles. Some observe. Some interpret. Some act. Others validate, coordinate, or learn.

A monitoring agent may continuously track sensor data and flag deviations. A diagnostic agent may correlate events and identify likely root causes. A planning agent may decide whether maintenance is required and when it should be scheduled. A connected worker agent may guide technicians through the correct procedure. A performance agent may analyze post-action metrics to confirm problem resolution. Together, these agents deliver comprehensive coverage across the entire operational cycle.

This specialization ensures agents are not overloaded with responsibilities but instead operate with clarity and precision. And because each agent is built on the same contextual intelligence provided by IRIS Foundry, they communicate using a shared understanding of equipment, processes, and workflows.

8.2 Reliability and Maintenance Agents

Reliability is one of the primary domains where industrial agents deliver transformative value. Reliability-focused agents monitor asset health, detect anomalies, and assess degradation patterns using real-time and historical data. They can interpret sensor drift, rising vibration signatures, temperature anomalies, lubrication patterns, and other early indicators of failure.

When an anomaly is detected, additional agents step in. A causal reasoning agent evaluates whether the anomaly aligns with known failure modes, maintenance history, or upstream disruptions. A maintenance planning agent decides whether the issue warrants immediate action, scheduled work, or additional investigation. A documentation agent retrieves relevant sections from manuals, P&IDs, or past inspection reports. Finally, an execution agent may initiate a work order, prepare an inspection plan, or escalate the issue to supervisors.

These reliability agents mimic the combined reasoning of analysts, planners, and technicians—ensuring potential failures are identified early and handled consistently.

8.3 Production and Process Optimization Agents

Production environments present a different set of challenges: fluctuating demand, process variation, quality deviations, and throughput constraints. IRIS Flows agents support production teams by

continuously evaluating process stability and performance.

A process monitoring agent keeps track of key variables such as pressure, temperature, flow rate, and energy consumption. A performance agent analyzes deviations to determine whether they impact yield, throughput, or cycle times. When inefficiencies arise, optimization agents evaluate operational alternatives, such as adjusting set points, modifying production schedules, or recommending preventive interventions.

In more advanced applications, orchestration agents coordinate multiple factors—maintenance schedules, raw material quality, operating limits—to propose optimized production configurations. These agents can help plants maintain stable, efficient, and predictable production outcomes even under rapidly changing conditions.

8.4 Safety, Compliance, and Environmental Agents

Safety is an area where variability in human behavior can introduce risk. IRIS Flows includes agents designed specifically to ensure compliance with safety procedures, environmental regulations, and operational policies.

Safety agents monitor for patterns such as equipment overheating, chemical imbalances, pressure spikes, or unsafe operating envelopes. When risks escalate, they initiate corrective workflows, notify supervisors, or recommend shutdown sequences. Compliance agents monitor procedural adherence, ensuring inspections are completed, documentation is up-to-date, and safety-critical tasks follow the correct sequence.

Environmental agents track emissions, waste streams, and energy usage, correlating operational behavior with environmental impact. These agents help organizations maintain safety and sustainability without constant manual oversight.

8.5 Alarm Management and Event Intelligence Agents

Alarm overload is a persistent challenge in industrial environments. Operators cannot manually triage dozens or hundreds of alarms during peak events, leading to missed signals and reactive firefighting.

IRIS Flows agents help by classifying, grouping, and prioritizing alarms based on real-time operational conditions. Instead of treating every alarm as isolated, agents analyze patterns to identify root causes behind alarm storms and recommend actions that address underlying issues rather than symptoms. These agents can also evaluate alarm histories, detect recurring patterns, and suggest changes to alarm configurations to reduce nuisance alerts. By automating triage and providing context, they reduce operator load and help teams focus on the events that matter most.

8.6 Connected Worker and Procedure Automation Agents

Through integration with Proceedix, IRIS Flows includes agents that bridge the digital and physical worlds. Connected worker agents generate dynamic, context-specific procedures and inspection workflows for technicians in the field. They pull relevant information from manuals, shift logs, and historical data, transforming it into step-by-step guidance delivered via mobile devices, tablets, or AR headsets.

These agents adapt instructions based on real-time conditions. For example, if a technician encounters a discrepancy during an inspection, the agent can modify the workflow, ask clarifying questions, or escalate the issue to a supervisor. This adaptability ensures technicians always follow the safest, most effective procedures—regardless of experience level.

These agents democratize expertise, enabling less experienced workers to perform complex tasks confidently and accurately.

8.7 Knowledge, Documentation, and Intelligence Retrieval Agents

Industrial operations depend heavily on documentation—but manuals, P&IDs, engineering drawings, inspection records, and corrective actions are often difficult to find and interpret under pressure.

IRIS Flows solves this with agents that specialize in knowledge retrieval and document intelligence. These agents can instantly extract relevant passages from long PDFs, identify the correct schematic for a piece of equipment, interpret complex diagrams, or summarize historical maintenance entries related to a current anomaly. They reduce the cognitive burden on operators and accelerate decision-making significantly.

Because these agents are grounded in the knowledge graph, they understand not only the content of the documents but also the assets and processes those documents influence.

8.8 Multi-Agent Collaboration: The Digital Workforce in Action

While each agent type is valuable independently, the greatest operational improvement occurs when agents work collaboratively. IRIS Flows enables multi-agent chains and ecosystems where specialized agents contribute their insights to solve complex problems.

A malfunctioning compressor, for instance, may trigger interactions among:

- Monitoring agents (detecting the anomaly)
- Diagnostic agents (identifying the cause)
- Optimization agents (assessing impact on production)
- Connected worker agents (guiding technicians)
- Verification agents (confirming resolution)
- Knowledge agents (surfacing relevant documentation)

This collaborative approach mirrors how human teams operate—but with the speed and consistency only AI can deliver.

8.9 The Agent Ecosystem as a Competitive Advantage

The diversity and depth of the agent ecosystem within IRIS Flows provides organizations with an architectural advantage that traditional automation, isolated AI models, or simple copilots cannot match. Instead of relying on centralized intelligence or individual workflows, IRIS Flows distributes intelligence across many specialized agents, each contributing expertise to a shared operational goal.

This decentralized model increases resilience, reduces bottlenecks, supports parallel problem-solving, and enables continuous learning across the entire operational landscape. It is the cornerstone of SymphonyAI's vision for autonomous industrial operations.

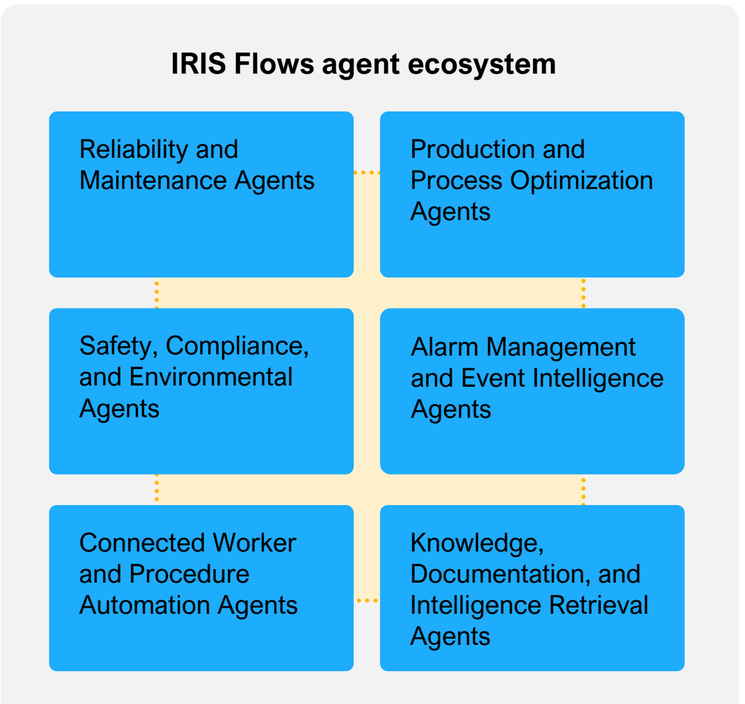

9. Use cases and real-world scenarios

IRIS Flows is not an abstract concept or a theoretical framework—it is a practical, deployable system already transforming how industrial operations respond to events, optimize processes, and elevate human performance. The following scenarios illustrate how IRIS Flows agents behave in real-world situations, demonstrating the power of agentic AI working in context, in coordination, and in real time.

These examples intentionally span reliability, maintenance, production, quality, safety, and connected worker domains to show the wide-ranging applicability of the IRIS Flows digital workforce.

9.1 Predictive → Prescriptive → Autonomous Maintenance

In a compressor station operating continuously under high load, a subtle vibration rise is detected—well below alarm thresholds but outside typical trend patterns. Under conventional systems, this deviation might go unnoticed for days or be flagged without any meaningful diagnosis, leaving the decision-making burden on overextended reliability engineers.

With IRIS Flows, a Monitoring Agent immediately observes the unusual pattern and correlates the signal with recent temperature fluctuations and slight pressure instability upstream.

It pulls historical records from the knowledge graph and identifies an event from two years earlier where a similar pattern preceded a bearing failure.

The Monitoring Agent hands this context to a Diagnostic Agent, which uses the Industrial LLM and asset-specific failure modes to evaluate the likely root cause. It determines that lubrication degradation is a probable cause and identifies the exact component most at risk.

A Maintenance Planning Agent then assesses the risk profile and compares it to current production priorities, workforce availability, and safety constraints. Recognizing that the failure could develop rapidly under current load, it recommends an inspection within the next two hours.

Through IRIS Flows, a Connected Worker Agent automatically generates a step-by-step inspection procedure customized for this compressor and sends it to the nearest qualified technician using Proceedix. Once the inspection is completed, a Verification Agent checks post-inspection metrics to determine whether the issue has been corrected or whether further action is needed.

This entire sequence—monitor, diagnose, plan, guide, and validate—occurs within minutes, without requiring engineers to sift through dashboards, logs, or maintenance systems manually.

9.2 Alarm Storm Management: From Chaos to Clarity

A petrochemical plant experiences a sudden surge of alarms triggered by a blocked heat exchanger. Traditional alarm systems merely flood the operator console with alerts—temperature deviations, pressure spikes, flow imbalances—forcing operators to prioritize under intense pressure.

In IRIS Flows, the moment the first alarm is detected, an Alarm Interpretation Agent begins analyzing the pattern. It recognizes that the signals are not independent but interconnected symptoms of a single issue. It clusters the alarms, identifies the likely root cause, and suppresses irrelevant secondary alerts.

A Causal Reasoning Agent correlates the event with upstream sensor data and recently completed maintenance activities, concluding that fouling is likely the cause. Meanwhile, a Production Impact Agent evaluates the downstream consequences: potential feed disruptions, quality degradation, and energy inefficiencies.

With this coordinated intelligence, a Decision Agent recommends a controlled slowdown rather than an emergency shutdown. It triggers a workflow to notify both control room operators and field technicians, all while maintaining a detailed log of reasoning for audit purposes.

The result is a calm, coordinated response instead of an overwhelming storm of alarms.

9.3 Quality Deviation Investigation and Rapid Root Cause Analysis

In a batch-based manufacturing environment, a deviation in product density is detected during routine sampling. Historically, such deviations required hours of manual investigation across batch records, operator logs, historian data, and equipment histories.

With IRIS Flows, a Quality Monitoring Agent immediately detects the deviation and retrieves the batch genealogy using the knowledge graph. It correlates the deviation with variations in raw material properties, slight temperature fluctuations, and a valve actuation delay mid-cycle.

A Document Intelligence Agent scans SOPs, prior deviation reports, and operator notes from the past six months.

It finds a recurring note about a sticky valve that occasionally causes delayed actuation.

A Diagnostic Agent evaluates this as the most likely cause and recommends both a targeted maintenance check and a short-term process adjustment to stabilize the next batch.

The system produces a comprehensive root cause report—what happened, why it happened, what to do now, and what to prevent next—within minutes instead of hours or days.

9.4 Autonomous Shift Handover and Operational Awareness

A common challenge in multi-shift operations is the loss of context during handovers. Critical details can be missed, events may go unnoticed, and tribal knowledge remains trapped in conversations rather than recorded systematically.

IRIS Flows introduces a Shift Intelligence Agent that continuously monitors operations during each shift—tracking anomalies, maintenance tasks, procedure deviations, and significant operational changes. Throughout the shift, the agent builds a structured, chronological narrative enriched with contextual cues from the knowledge graph.

At handover time, the agent generates a clear, concise, insight-rich shift summary. It highlights anomalies that require attention, ongoing maintenance actions, deviations from normal patterns, and any unusual operator interventions. The incoming shift receives a fully contextualized operational snapshot rather than fragmented information.

This ensures continuity, reduces errors, and strengthens situational awareness across all shifts.

9.5 Energy Optimization Through Multi-Agent Coordination

In an energy-intensive process, small inefficiencies can lead to significant cost increases. IRIS Flows uses a team of agents to optimize energy use continuously.

A Monitoring Agent tracks energy consumption patterns, while an Optimization Agent correlates them with production schedules, equipment efficiency curves, and live process conditions. A Constraints Agent ensures no optimization recommendations violate safety limits or product quality requirements.

When inefficiencies arise—such as a pump running above its optimal energy curve—the Optimization Agent recommends setpoint adjustments or alternate machine sequencing. If approved, an Execution Agent performs the adjustments or alerts operators through existing control systems.

These multi-agent optimizations can reduce energy consumption by 5–15% without compromising throughput or quality.

9.6 Connected Worker Execution for Safety-Critical Tasks

During an inspection in a hazardous environment, a field technician encounters a condition not covered in the standard procedure. Instead of guessing or returning to the control room, the technician communicates with a Connected Worker Agent.

This agent immediately retrieves relevant procedures, references historical inspection notes, and analyzes real-time sensor data related to the asset. It identifies three scenarios where similar anomalies were observed and suggests the safest corrective action. If the scenario requires supervisor approval, the agent notifies the appropriate person instantly.

Through AR-enabled guidance, the agent walks the technician step-by-step through the correct action. Once the task is completed, a Reporting Agent updates relevant systems automatically, ensuring documentation is complete and compliant.

This creates a strong bridge between AI intelligence and frontline execution.

9.7 Production Yield Improvement Through Intelligent Collaboration

A continuous manufacturing line struggles with fluctuating yield due to upstream flow variability. In conventional operations, this would require repeated manual tuning and after-the-fact analysis.

With IRIS Flows, a Process Stability Agent monitors flow patterns at high frequency, flagging deviations that correlate with yield drops. A Correlation Agent identifies which upstream conditions most strongly influence the variations, while a Predictive Agent models the likely impact of different corrective actions.

An Optimization Agent then proposes a series of adjustments—flow tuning, temperature balancing, and modified cycle durations—validated against process constraints. After human confirmation or automated approval, an Execution Agent implements these changes, and a Verification Agent monitors the short- and long-term effects to update the knowledge base.

This closed-loop learning improves yield consistently across shifts and operating conditions.

10. Business impact and ROI

Industrial organizations face intense pressure to increase efficiency, reduce downtime, maintain safe operations, and achieve more with fewer resources. Yet many continue to rely on manual processes, fragmented systems, and overburdened teams who must interpret complex signals and make rapid decisions under stress. IRIS Flows directly addresses these challenges by introducing an intelligent digital workforce—industrial agents that monitor, analyze, decide, and act. The result is not only improved operational performance but also measurable, enterprise-wide business value.

IRIS Flows delivers ROI through five primary levers: reduced downtime, improved maintenance efficiency, enhanced workforce productivity, accelerated decision-making, and enterprise-scale consistency. These levers compound across assets, shifts, and sites, creating transformational impact rather than incremental gains.

10.1 Reducing Unplanned Downtime

Unplanned downtime remains one of the largest hidden costs in industrial operations. When equipment fails unexpectedly, production stops, maintenance expenses spike, schedules are disrupted, and downstream quality issues multiply.

Traditional predictive maintenance solutions can identify early warning signs, but they often fail to bridge the gap between detection and timely human response.

IRIS Flows closes this gap by automating the entire chain from anomaly detection to diagnosis, action, and validation. Monitoring agents interpret subtle deviations long before alarms occur; diagnostic agents identify likely causes using knowledge graph correlations; maintenance agents determine the appropriate intervention; and connected worker agents deliver guided procedures to technicians.

By ensuring early detection and consistent action, IRIS Flows significantly reduces the frequency and severity of unplanned failures. Even a modest reduction in unplanned downtime—5% to 10%—can yield millions of dollars in savings for large facilities. For organizations with highly interdependent operations or expensive assets, the impact can be even more dramatic.

10.2 Increasing Maintenance Efficiency and Accuracy

Most maintenance organizations struggle with reactive work, variable execution quality, and incomplete documentation.

Even when maintenance is planned, the complexity of modern equipment means technicians often spend excessive time searching for the correct procedures, interpreting historical notes, or determining the right steps to take.

IRIS Flows transforms maintenance execution by guiding technicians with context-aware, step-by-step workflows generated in real time. Agents retrieve the correct procedures from documents, tailor instructions to current asset conditions, and update CMMS systems automatically. This reduces the variability that often leads to recurring failures or quality issues, while standardizing high-quality execution across the workforce.

Maintenance teams also benefit from optimized scheduling, proactive inspections, and improved prioritization of tasks based on actual risk rather than outdated heuristics. These improvements translate directly into lower maintenance costs, longer asset life, and fewer repeated interventions.

10.3 Amplifying Workforce Productivity Across All Roles

Industrial organizations today face significant workforce challenges: aging expertise, high turnover, slow onboarding, and increasing technical complexity.

IRIS Flows does not replace workers; it elevates them. Agents serve as digital teammates that augment human capabilities, reduce cognitive load, and provide expert-level guidance on demand.

Operators gain access to real-time insights, automated prioritization, and clear decision support. Engineers receive automatically generated diagnostics, historical correlations, and structured investigations. Technicians benefit from AR-guided workflows and instant access to the precise steps needed to resolve issues. Supervisors get automated shift summaries, performance insights, and detailed visibility into plant-wide activity.

This augmentation not only improves productivity but also reduces errors, speeds up problem resolution, and equips less experienced personnel to perform at higher levels. Organizations can close skill gaps without relying solely on hiring or long training cycles, creating a more resilient and capable workforce.

10.4 Accelerating Decision-Making and Response Times

Every industrial facility encounters moments where fast, accurate decision-making is essential. Alarm storms, quality deviations, equipment anomalies, and unexpected changes in feedstock or production conditions require immediate attention. Delayed or inconsistent decisions often lead to cascading failures.

With IRIS Flows, decision-making becomes both faster and more reliable. Agents monitor data continuously, correlate signals automatically, and initiate action sequences without waiting for humans to piece together information from disparate systems. When human intervention is needed, agents present the actionable context—what happened, why it matters, what the probable causes are, and what should be done next.

This speed and clarity dramatically reduce the time between issue detection and corrective action. In many cases, intervention happens before operators are even aware something is amiss. Faster decisions translate to fewer disruptions, more stable operations, and higher output.

10.5 Creating Enterprise-Scale Consistency and Standardization

Large industrial organizations often struggle with variability across sites. Different operating cultures, tooling, documentation, and personnel lead to inconsistent results. Even when best practices exist, they may not be consistently applied across shifts or locations.

IRIS Flows standardizes excellence at scale. Agents execute tasks the same way every time, guided by enterprise-defined rules, procedures, and operational models. When new insights or best practices emerge, they can be embedded into the agent ecosystem and instantly propagated across every site. This ensures global alignment in how anomalies are handled, how procedures are followed, and how decisions are made.

The result is a more predictable, reliable, and scalable operational model that reduces risk and elevates performance across the entire organization.

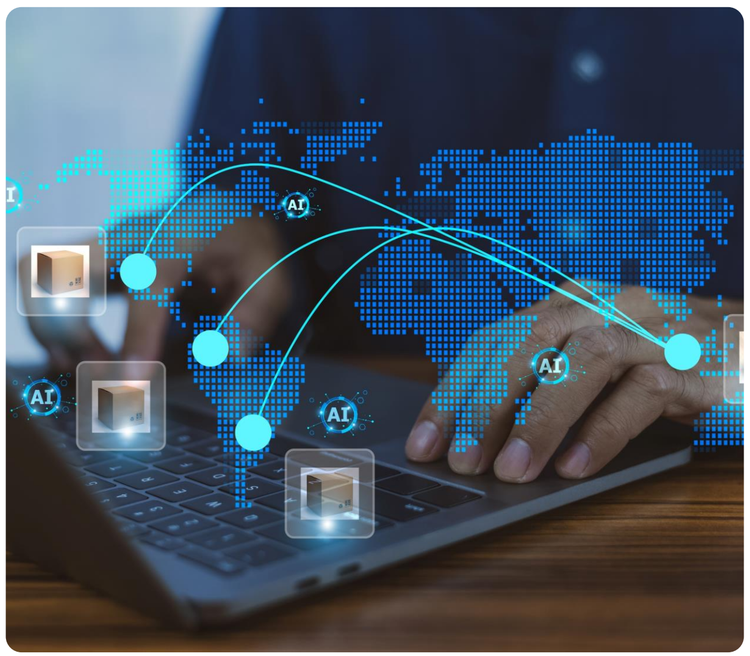

11. Differentiation

IRIS Flows differentiates itself in five structural ways that, together, make it the most capable agentic AI environment for asset-intensive industries.

11.1 Built on a Semantic Industrial Knowledge Graph

At the heart of IRIS Flows is the semantic industrial knowledge graph inside IRIS Foundry. This graph models assets, processes, signals, documents, equipment hierarchies, engineering relationships, and historical patterns in a way that mirrors how plants actually operate. Unlike generic AI systems that treat data as isolated text or numbers, IRIS Flows agents interpret events within a precise, interconnected operational model.

When an agent detects a deviation, it does not analyze the signal in isolation. It understands how that signal relates to upstream pressure, downstream quality, maintenance actions performed six months earlier, and the equipment hierarchy mapped in engineering drawings. It understands the asset's operating envelope, typical failure patterns, and surrounding process context. This grounding prevents hallucinations, ensures accurate interpretation, and gives agents the factual foundation needed for safe autonomy.

No other industrial AI platform offers this depth of contextual intelligence—and without it, agentic AI in industry cannot be trusted to operate safely or effectively.

11.2 Powered by a Domain-Trained Industrial LLM

IRIS Flows draws its reasoning from SymphonyAI's Industrial LLM, a large language model trained on an extensive corpus of engineering knowledge, operational terminology, maintenance documentation, P&IDs, FMEAs, troubleshooting guides, shift logs, and real-world industrial data. This training gives IRIS agents domain fluency—an understanding of industrial vocabulary, equipment behavior, physics-informed logic, and failure relationships that generic LLMs simply cannot replicate.

The Industrial LLM is deeply integrated with the knowledge graph, enabling dual reasoning: probabilistic language understanding combined with deterministic contextual grounding. As a result, agents can interpret technical documentation, summarize engineering insights, explain root causes, and guide workers with domain-specific clarity.

Generic LLMs, even when fine-tuned, struggle to understand the intricacies of industrial processes or apply safe decision-making rules. IRIS Flows avoids these pitfalls by building intelligence directly from industrial reality rather than adapting general-purpose models.

11.3 Enabled by a Composable Industrial DataOps Foundation

Industrial AI is only as good as the data it depends on. Many AI systems fail because they rely on static, uncontextualized, or incomplete data. IRIS Foundry's Industrial DataOps architecture solves this by providing a real-time, composable data backbone that harmonizes data from IT, OT, IoT, engineering repositories, and external enterprise systems.

This DataOps layer ensures that agents operate on high-quality, timely, trusted data—data that is contextualized, validated, and mapped into the knowledge graph automatically. It manages ingestion of millions of data points per second, maintains lineage across systems, and normalizes information into a unified namespace that agents use for consistent reasoning.

This makes IRIS Flows fundamentally different from AI systems that rely on ad-hoc integrations or brittle pipelines. Its DataOps foundation ensures reliability, scalability, and enterprise-grade governance, enabling agents to perform with precision at all hours, under all conditions.

11.4 Integrated with IRIS Copilot and Proceedix for Closed-Loop Execution

One of the biggest weaknesses of traditional AI systems is the “last mile” problem: insights do not lead reliably to action. IRIS Flows overcomes this by integrating seamlessly with IRIS Copilot and Proceedix Connected Worker, creating a tightly coupled loop from detection to action to validation.

IRIS Copilot serves as the conversational interface through which operators, technicians, and engineers can query agents, request explanations, refine decisions, or escalate issues. Proceedix extends the loop into the physical world, delivering dynamic, context-specific, and often AR-enhanced workflows directly to frontline workers.

Agents in IRIS Flows do not stop at recommending actions—they initiate work orders, deliver procedures, update systems, and verify outcomes. This closed-loop capability ensures operational continuity and eliminates the common disconnect between insights and execution.

By unifying reasoning and action, IRIS Flows bridges the digital and physical worlds in a way that no traditional analytics or AI tool can accomplish alone.

11.5 Designed for Enterprise Scale, Safety, and Industrial Governance

Industrial operations demand systems that are safe, resilient, and governable. IRIS Flows was built for this from the beginning, with comprehensive safety, governance, and control layers that define how agents behave, what actions they are allowed to take, and how decisions are monitored.

Every agent action is logged and auditable. Every autonomous step is bounded by predefined rules. Every workflow includes built-in safety constraints. Human override is always available, and agents must provide transparent justifications for their recommendations. These guardrails allow organizations to adopt autonomy gradually and confidently.

Equally important, IRIS Flows is designed to scale across thousands of agents operating simultaneously across multiple sites, regions, and infrastructures. Load balancing, distributed processing, multi-site synchronization, and high-availability mechanisms ensure that IRIS Flows remains responsive and consistent regardless of deployment size.

This enterprise-readiness is a critical differentiator. Experimental agent systems may work in isolated pilots, but only IRIS Flows can power continuous industrial operations across global networks.

11.6 A Transformational Advantage, Not an Incremental One

Taken together, these differentiators—contextual grounding, domain intelligence, data integrity, closed-loop execution, and enterprise governance—form a structural advantage that competitors cannot replicate simply by adding LLMs or workflow engines.

IRIS Flows is not a chatbot, not a scripting platform, not a dashboard, and not an RPA tool. It is a comprehensive, industrial-grade agentic AI environment that transforms operations by enabling teams of intelligent agents to monitor, analyze, recommend, execute, and learn—continuously, coherently, and safely.

This is why IRIS Flows stands alone. It is not an incremental improvement over existing systems; it is a fundamentally new operational paradigm, built specifically for the complexities and stakes of industrial environments.

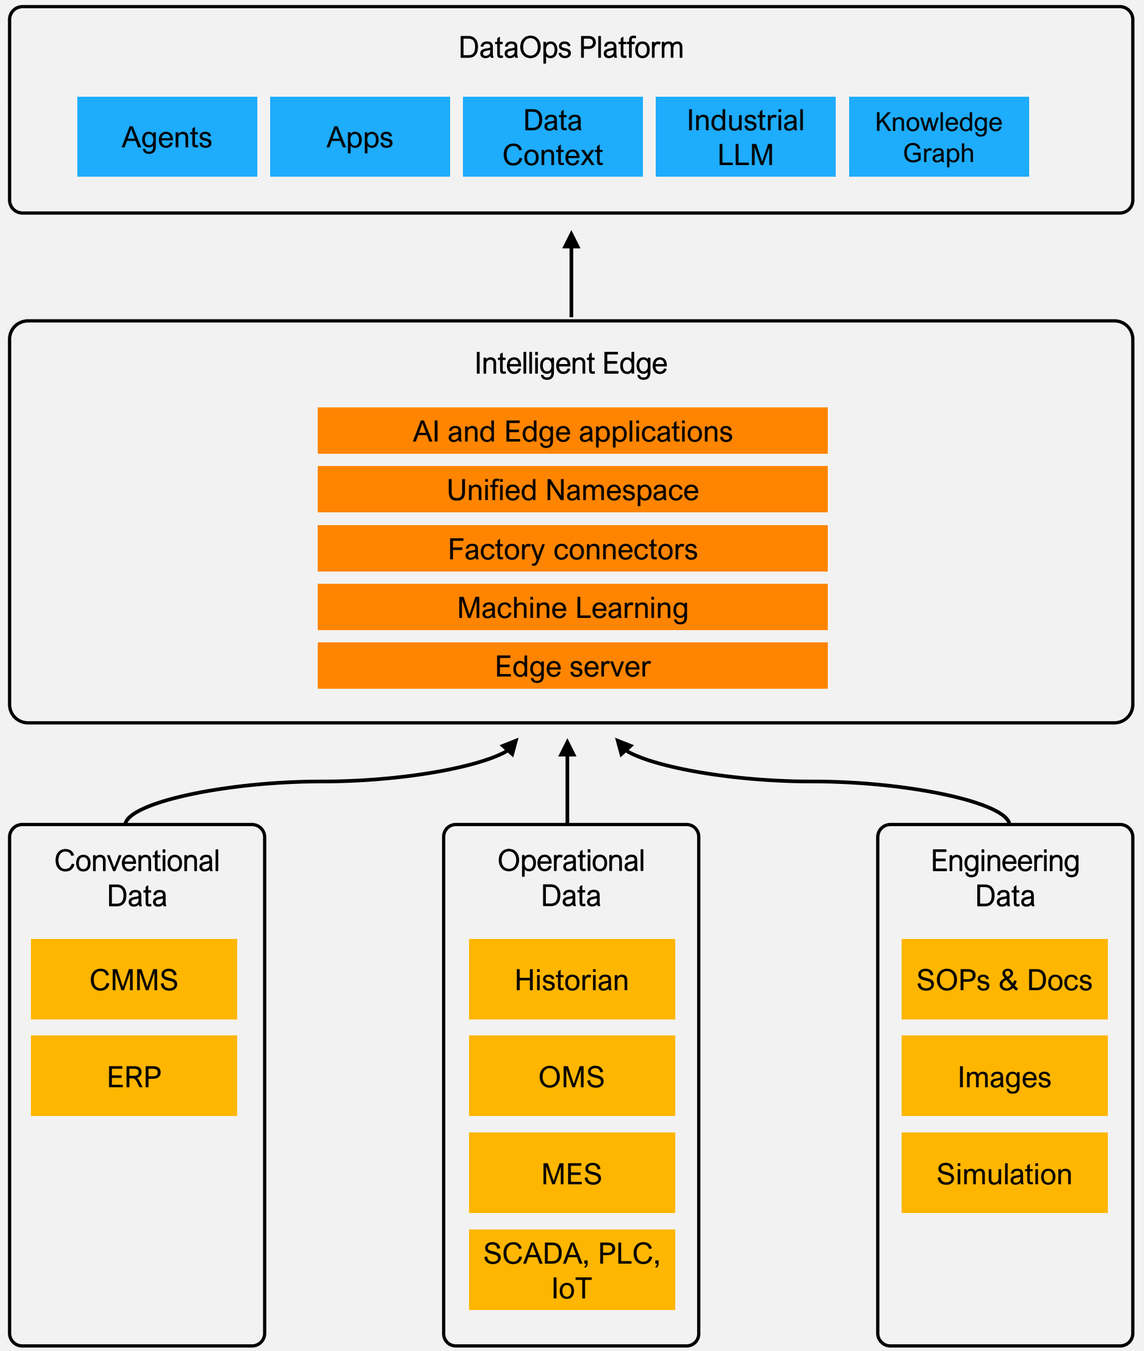

- Built on a Semantic Industrial Knowledge Graph
- Powered by a Domain-Trained Industrial LLM
- Enabled by a Composable Industrial DataOps Foundation
- Integrated with IRIS Copilot and Proceedix for Closed-Loop Execution
- Designed for Enterprise Scale, Safety, and Industrial Governance
- A Transformational Advantage, Not an Incremental One

12. The future of Agentic AI

Industrial operations have reached an inflection point. Digital transformation has delivered connectivity, data availability, and pockets of analytics-driven improvement—but the core operational model has largely remained the same: humans interpret data, decide what to do, and manually execute actions. As complexity rises and the workforce becomes increasingly strained, this model can no longer scale.

Agentic AI introduces an entirely new operational paradigm—one in which intelligence is distributed, autonomous, coordinated, and continuously learning. IRIS Flows embodies this paradigm, enabling organizations to augment human capabilities and gradually transition toward self-optimizing, adaptive operations. Yet the impact of agentic AI is only beginning. Over the next decade, industrial agents will evolve from decision assistants to operational collaborators and, ultimately, to autonomous stewards of mission-critical systems.

12.1 From Assistive Systems to Autonomous Operational Partners

Today, agentic AI primarily functions in assistive and augmented modes. Agents detect anomalies, perform diagnostics, recommend actions, and generate procedures—always with human oversight or final approval. This phased approach is essential for building trust and ensuring operational safety.

But as agents accumulate experience and learn from outcomes, organizations will gain the confidence to increase autonomy levels, allowing agents to handle low-risk interventions independently and to coordinate multi-step workflows with minimal human involvement. Over time, agents will become trusted operational partners capable of executing well-defined tasks autonomously while keeping humans informed through transparent reasoning and audit trails.

This trajectory mirrors the journey of automation in aviation and automotive industries: first assistive, then supervisory, then increasingly autonomous.

12.2 The Emergence of Fully Coordinated Multi-Agent Operations

As organizations deploy more agents across reliability, maintenance, quality, production, and safety functions, the true power of agentic AI will emerge in the form of multi-agent coordination at plant-wide and even fleet-wide scale.

Rather than relying on centralized intelligence, future operations will be governed by networks of specialized agents working in concert. These digital teams will continuously re-evaluate conditions, collaborate to resolve issues, and optimize performance across systems. One agent may monitor equipment health, another may manage throughput, while a third may ensure energy efficiency or compliance with environmental regulations. Together, they will form a dynamic, decentralized intelligence layer that continuously adapts to changing conditions.

In this future, the plant becomes a living system—autonomous, interconnected, and capable of self-correction.

12.3 Agents That Learn Not Just Individually, but Collectively

The evolution of industrial agents will be defined not only by individual learning but by the emergence of shared, collective learning. Every intervention, failure mode, corrective action, and performance improvement identified by an agent in one facility will contribute to the knowledge base used across all facilities.

A bearing failure resolved in one plant will inform agents monitoring similar assets elsewhere. A setpoint optimization discovered on one production line will be shared across global operations. A recurring safety deviation in one region will become a proactive alert in another.

This collective learning transforms every plant incident into enterprise-wide intelligence, accelerating continuous improvement far beyond what human teams could achieve alone.

12.4 Digital Twins Will Evolve into Agentic Twins

Traditional digital twins replicate physical systems but remain primarily passive—visualization tools that require human interpretation. Agentic AI will elevate digital twins into agentic twins: active, continuously reasoning digital environments where agents run simulations, test scenarios, validate decisions, and anticipate operational outcomes before acting.

These agentic twins will allow organizations to preview the impact of maintenance actions, production changes, or environmental shifts. They will run continuous "what-if" analyses in the background, ensuring operations are optimized, stable, and resilient.

The plant of the future will be guided not only by real-time data, but by predictive simulations executed autonomously by agentic twins.

12.5 A Workforce Empowered by Intelligent Companions

As agentic AI matures, it will fundamentally reshape the dynamics of the industrial workforce. Human expertise will remain essential, but the burden of constant monitoring, manual analysis, procedural recall, and cross-system coordination will diminish significantly.

Operators will work alongside intelligent companions capable of providing real-time insight and guidance. Technicians will perform tasks with confidence, supported by dynamic procedures generated by agents that understand each asset's unique history and condition. Engineers will focus on innovation and optimization rather than firefighting. Supervisors will oversee operations with greater clarity and foresight.

Organizations that embrace this augmented workforce model will be positioned to thrive in the face of labor shortages, increasing complexity, and rising performance expectations.

Preparing for the Agentic Future

The organizations that will excel in the next decade will be those that begin adopting agentic AI today. The path forward begins with augmenting human teams with assistive agents, evolving to coordinated multi-agent ecosystems, and ultimately embracing autonomous capabilities supported by agentic twins and continuous learning systems.

IRIS Flows offers a structured, safe, and scalable way to take this journey. It is not simply a product; it is the foundation for the autonomous industrial enterprise of the future.

A Foundation for Long-Term Transformation

The value created by IRIS Flows extends far beyond immediate cost savings. By shifting from reactive operations to predictive, prescriptive, and increasingly autonomous modes, organizations build a foundation for long-term competitiveness. Agentic AI becomes a strategic differentiator—one that enables plants to operate with consistency, intelligence, and agility even as workforce dynamics, regulatory landscapes, and production demands evolve.

IRIS Flows does not simply optimize existing processes; it transforms the operational model itself. This positions organizations to unlock new efficiencies, new capabilities, and new levels of reliability that traditional automation and analytics could never achieve.

SymphonyAI

Why SymphonyAI Industrial

Learn more about IRIS Foundry

Learn more about IRIS Flows

IRIS Flows Demo

Contact Sales