


E-BOOK

Why Enterprise AI Stalls

And the missing foundation required to scale it





Contents

Introduction: The AI Value Gap	3
Chapter 1: Why Enterprise AI Struggles to Scale	5
Chapter 2: The Foundation for Scalable Enterprise AI	8
Chapter 3: Focus Investment Where It Matters	11
Chapter 4: Equip AI to Operate Reliably	14
Chapter 5: Make AI Performance Visible	17
Chapter 6: Creating a Continuous AI Improvement Model	20
Conclusion	23

Introduction:

The AI Value Gap

AI capability is advancing rapidly but the leadership challenge is converting that capability into measurable, repeatable enterprise value.

Enterprise AI has entered a new phase. The first phase was about proving that the technology could work. Organizations experimented with generative AI, deployed assistants, tested agents and automated individual tasks. Those initiatives demonstrated genuine capability and created significant momentum.

The question facing leaders today is different—**How do we convert AI capability into reliable and cost-effective operations across the enterprise.**

There is no doubt, AI investment IS increasing. The number of models, platforms, agents and frameworks continues to grow. Yet many organizations still struggle to translate that trend into sustained improvements in productivity, cost, customer experience, control or growth.

Technology alone does not guarantee innovation, and successful pilots do not automatically become successful enterprise programs. As we have seen time and again, what works well in a controlled lab, often does not scale.

As AI moves into real operations, leaders face an increasingly complex set of decisions:

- ❑ Where should the organization invest?
- ❑ Which processes and use cases offer meaningful value?
- ❑ How should AI operate within existing controls and responsibilities?
- ❑ How can agents act consistently across end-to-end business processes?
- ❑ How should AI performance be measured?
- ❑ How can context, governance and operating assumptions remain current as the business changes?
- ❑ How can successful deployments be expanded without increasing cost and risk at the same rate?

These are not simply technology questions. They are questions about how the enterprise operates.

Models demonstrate broad general knowledge, and through generic skills can generate, predict, recommend and act. Yet despite this vast knowledge they do not inherently understand how a specific organization is designed to work. They do not automatically know its operating rules, process dependencies, authority limits, control requirements or exception paths.

Without that understanding, AI can perform tasks while still failing to improve the wider business outcome. Furthermore, without adequate guidance the agent is left to determine the best approach, and reasoning through a process will consume large quantities of tokens increasing the operating cost while delivering little value-add. The organizations that scale AI successfully will therefore do more than deploy increasingly capable technology. They will build an operating foundation that allows them to focus investment, equip AI to operate reliably, connect AI activity to business performance and improve continuously.

That foundation is Process Context.


Process Context connects AI to the reality of the enterprise. It gives people and AI a shared understanding of how work should happen, how it actually happens and what must change to improve the outcome. The result of this is agents that are better directed, more reliable, more measurable and better equipped to create enterprise value.

Put simply, would you trust a surgeon who has never spoken to you or looked at your chart? Probably not.

Think of AI like an incredibly skilled surgeon. However, even the best surgeon cannot operate safely by looking at a patient for the first time in the operating room. They need the patient's medical history, test results, treatment plan, and knowledge of what has happened before.

Process Context plays the same role for AI. It provides the history, procedures, rules, and real-world operating conditions of the enterprise. The AI doesn't become smarter. The AI doesn't become smarter—it becomes better informed, making better decisions, following the right processes, measuring results and delivering better business outcomes.

Task, data, knowledge, operational context—your AI stack already has all four covered. What it doesn't have is a fifth: process context—how your business actually runs.



Chapter 1:

Why Enterprise AI Struggles to Scale

The transition from a contained pilot to enterprise operations exposes gaps that stronger models and more orchestration cannot resolve on their own.

The transition from experimentation to enterprise deployment exposes challenges that are easy to overlook during a pilot. A pilot is intentionally contained—it usually has a defined task, a limited population, accessible data and a team focused on making the initiative succeed.

Enterprise operations are different. Work crosses functions, systems, locations and organizational, and sometime even geographical boundaries. Decisions depend on policies, responsibilities and business conditions. Exceptions are common. Outcomes are shaped by a chain of events rather than a single task.

An AI agent operating in this environment needs more than a strong model and access to data. It needs to understand the operating context around the task.

Many AI initiatives begin with a narrow objective:

- ☐ Generate a response
- ☐ Classify a request
- ☐ Recommend an action
- ☐ Update a record
- ☐ Approve a transaction
- ☐ Resolve an exception
- ☐ Trigger a workflow

The agent may complete that task successfully. However, task completion does not demonstrate that the end-to-end process has improved.

An action can be technically correct while producing an undesirable operational consequence. An agent could accelerate one stage of a process while creating additional work downstream. It could follow a general instruction while overlooking an enterprise-specific policy. Or it could improve local productivity while increasing exceptions, control failures or customer delays elsewhere.

Key distinction

AI output and business outcome are not the same. The utility of an AI action is determined by whether it improves the performance of the process in which it operates, while at the same time adhering to our compliance and regulatory controls. For example, don't just onboard customers faster. Onboard them faster while maintaining compliance, risk controls, and customer quality.

Common enterprise scaling barriers:

Fragmented investment decisions

Use cases are often initiated by individual functions, technology teams or enthusiastic sponsors. This creates momentum, but it can also produce a fragmented portfolio with inconsistent business cases and overlapping capabilities. With AI playing a major role in so many of their company's strategies, leaders may not know how many AI initiatives are actually under way, nor have a clear view of which initiatives will materially improve performance.

Incomplete operational understanding

AI teams can access documents, data and applications without having a connected understanding of the work itself. Important knowledge remains distributed across process models, operating procedures, policies, systems and experienced employees. As a result, each new initiative must reconstruct the context it needs.

Task-level optimization

Teams identify activities that can be automated but do not always consider their role within the end-to-end process. This can shift effort or delay rather than remove it. Local improvements can create new downstream constraints, exceptions or control requirements.

Governance, controls and auditability separated from execution

AI governance is frequently defined at a programme or policy level, while controls are documented elsewhere and audit evidence is gathered after the event. The challenge is to translate controls into enforceable process-level rules: the limits on what AI may do, the approvals and segregation of duties required, the evidence that must be retained, and the exceptions that must be escalated. Without this connection, organizations may be able to state how AI should behave but cannot reliably demonstrate that controls operated as intended or reconstruct why a particular action was taken.

Limited performance visibility

Many organizations can measure AI usage and outputs. Far fewer can demonstrate how AI activity affects end-to-end operational performance. Without that connection, it becomes difficult to validate value, diagnose underperformance or decide where further investment is justified.

Context that becomes outdated

Processes, systems, responsibilities and regulations change. If the context provided to AI does not change with them, reliability declines over time.

These challenges are connected. Together, they reveal a gap between what AI can do and what the enterprise needs AI to understand.

The missing operational understanding

Enterprise AI stacks commonly provide several forms of context:

- ❑ Task context: what the AI has been asked to do
- ❑ Data context: the information available to complete the task
- ❑ Knowledge context: documents, policies and reference material
- ❑ Operational context: the current state of an application, transaction or workflow

All are important but what is frequently missing is Process Context.

Process Context explains:

- ❑ How work fits into an end-to-end process
- ❑ What should happen before and after an activity
- ❑ Which dependencies affect the outcome
- ❑ Which roles have authority to act
- ❑ Which policies and controls apply
- ❑ Which conditions create an exception
- ❑ When an issue should be escalated
- ❑ How an action affects downstream performance
- ❑ How actual execution compares with intended execution

This is not supplementary information—it is the operating logic of the enterprise.

Why adding more AI does not resolve the problem

When results fall below expectations, the natural response is often to strengthen the technology stack. Organizations add more orchestration, more specialized agents, more knowledge sources or more detailed prompts.

These investments may improve individual capabilities. They do not automatically provide an understanding of how the enterprise operates. In simple terms: If five people are trying to cook dinner without the recipe, adding ten more people won't fix the problem. It just creates a bigger, more expensive mess.

Deploying additional AI without addressing this gap can increase the number of disconnected use cases, the cost of reproducing context, the volume of exceptions, the difficulty of measuring value and the effort required for governance.

Chapter takeaway:

The problem is not a lack of AI capability. It is a lack of connected enterprise understanding. Scaling successfully requires a foundation that closes that gap.

Chapter 2:

The Foundation for Scalable Enterprise AI

AI needs a reliable representation of the business it is expected to support, connecting operating intent with evidence of real execution.

Enterprise AI needs a reliable representation of the business it is expected to support. However, that representation must go beyond static documentation or isolated system data. It must connect how the organization intends work to happen with evidence of how work actually happens. This creates a shared operational foundation for people, transformation programmes and AI.

Connecting intent with operational reality

Every enterprise has at least two views of its operations:

The designed view

The designed view describes how the organization intends work to happen. It includes:

- Processes and activities
- Roles and responsibilities
- Standard operating procedures
- Policies and controls
- Decision logic
- Regulatory requirements
- Systems and information
- Escalation paths
- Target performance


The observed view

The observed view is derived from operational execution. It includes:

- System events
- Business objects
- Process variations
- Cycle times
- Bottlenecks
- Exceptions
- Rework
- Control deviations
- Actual performance

Both views are necessary. The designed view without execution evidence can become disconnected from reality. The observed view without design lacks the intent, rules and responsibilities needed to interpret what is happening.

AI needs the relationship between the two. It needs to understand what should happen, what is happening and what that means for the action it is expected to take.



A Digital Twin of how the enterprise operates

A Digital Twin creates a connected representation of enterprise operations. In this context, the Digital Twin is not merely a process diagram or a set of performance dashboards. It is a living foundation that connects business design, operational knowledge and execution evidence.

It provides a structured understanding of:

- How work flows across the enterprise
- Which people and systems participate
- Which rules and controls apply
- How process elements depend on one another
- Where actual execution differs from intended execution
- How those differences affect performance
- Where improvement or automation could create value
- How performance is measured

This is the foundation from which trusted Process Context can be created. For AI leaders, the importance is practical. The Digital Twin can provide a common operational reference for identifying opportunities, specifying expected behavior, supplying context to connected AI and evaluating performance after deployment.

It reduces the need for each AI initiative to reconstruct the enterprise from fragmented documents, interviews and system definitions.

One foundation, two enterprise outcomes

Many organizations continue to treat operational improvement and AI as separate programmes. One team is responsible for process transformation. Another is responsible for AI. Each has different technologies, budgets, governance and measures of success.

This separation creates duplication. AI teams need to understand how work happens in order to deploy effective use cases. Process and transformation teams already hold much of that knowledge. Process teams need better ways to turn insight into action. AI can help accelerate that action. A shared operational foundation connects the two agendas.

Outcome one:

Improved operational performance

- Understand end-to-end execution
- Identify bottlenecks and root causes
- Compare actual and intended performance
- Evaluate transformation opportunities
- Coordinate improvement across organizational boundaries
- Maintain a common view of processes, controls and responsibilities

Outcome two:

Scalable enterprise AI

- Identify high-value AI opportunities
- Give AI access to approved operational context
- Define rules, boundaries and escalation conditions
- Connect AI activity to end-to-end performance
- Monitor the effect of AI within business execution
- Improve future deployments based on evidence

The organization does not need to create a separate representation of the enterprise for every new agent or AI initiative. It can establish one governed foundation that supports multiple business and technology outcomes.

Process Context as reusable enterprise infrastructure

Without a shared foundation, context is repeatedly recreated. AI teams review documents, interview subject-matter experts, analyze systems and define process assumptions for each use case. The same operating knowledge is translated again for the next agent, platform or function. This increases cost and creates inconsistency.

A reusable Process Context foundation changes the economics of scale. Knowledge about processes, controls, roles and dependencies can be structured once, governed and made available across multiple initiatives.

The benefit is not only faster deployment. It is greater consistency between deployments. Different agents can draw from the same approved understanding of how the enterprise works. Changes to that understanding can be governed centrally rather than reproduced across an expanding portfolio of disconnected implementations.

Chapter takeaway:

Process Context should be seen as part of enterprise infrastructure rather than solely project documentation. The same foundations supports operational improvement and AI at scale which documentation alone cannot do.

CASE STUDY

ARIS and Boots

As part of a major SAP S/4HANA transformation, leading UK retailer Boots used ARIS to establish an enterprise-wide BPM capability and create a connected view of its processes, data and operations. Today, more than 2,000 detailed processes are documented, creating a trusted operational foundation for continuous improvement and future AI initiatives.

That visibility is already delivering measurable value. One finance process was redesigned from 220 steps to just 40, reducing execution time from six hours to 1.5 hours—a 75% improvement.

Boots is now building on this foundation by exploring Process Mining, Simulation, workflow orchestration and AI-powered assistants. Crucially, it is strengthening process governance and data quality before scaling AI.

The lesson is clear: before AI can transform how work gets done, it needs a trusted understanding of the business it is being asked to transform.



“If we don’t prepare our foundations, our data and our processes, that’s a fundamental mistake.”

Lee Oates

Head of Finance BPM & Continuous Improvement,
Boots

Chapter 3:

Focus Investment Where It Matters

The ability to identify where AI should be applied is as important as the ability to deploy it.

Enterprise AI creates almost unlimited possibilities. Leadership capacity and investment are not unlimited. The ability to identify where AI should be applied is therefore as important as the ability to deploy it.

The use-case portfolio problem

Many AI portfolios grow organically. A function identifies a problem. A sponsor secures funding. A technology team builds a pilot. Another function launches a similar initiative using a different platform. Individually, each initiative may appear reasonable. Collectively, the portfolio can become difficult to govern and even harder to prioritize.

Typical symptoms include:

- ❑ Large numbers of pilots with limited production adoption
- ❑ Similar use cases being developed independently
- ❑ Business cases based on activity rather than measurable outcomes
- ❑ Investment concentrated on what is easiest to automate
- ❑ Use cases proceeding without understanding downstream effects
- ❑ Pilots running with no clear understanding of the KPIs influenced
- ❑ Limited evidence for deciding which initiatives should scale

AI leaders need a more disciplined approach and a better starting question. Do not begin with “Where could AI be used?” Begin with “Where can AI create a sufficiently valuable, feasible and governable improvement?”

Start with the business outcome

Strong AI opportunities begin with a clearly defined business need.

That may include:

- ❑ Reducing end-to-end cycle time
- ❑ Lowering the cost of handling exceptions
- ❑ Improving customer response or resolution
- ❑ Increasing process capacity
- ❑ Reducing rework
- ❑ Improving compliance
- ❑ Accelerating a critical decision
- ❑ Reducing operational risk
- ❑ Improving working capital
- ❑ Increasing revenue conversion

The value case should be explicit before the technology choice is made. This anchors the initiative in an outcome the business recognizes and gives leadership a basis for comparing different investment options.

Assess the process, not only the task

A task may appear well suited to AI while the wider process is not ready for it.

For example, the activity may depend on poor-quality inputs, undocumented decisions, inconsistent hand-offs or policies that vary by market. Automating the activity without addressing those conditions may simply accelerate instability.

Effective prioritization examines factors such as:

- ❑ Process frequency and volume
- ❑ Current cost and performance
- ❑ Level of variation
- ❑ Frequency and type of exceptions
- ❑ Availability and quality of operational context
- ❑ Dependence on judgement
- ❑ Control and compliance requirements
- ❑ Consequences of an incorrect action
- ❑ Integration requirements
- ❑ Potential effect on downstream activities
- ❑ Expected financial or operational value

This allows leaders to distinguish between opportunities suited to AI, opportunities better served by conventional automation, opportunities requiring process redesign first and opportunities where the potential value does not justify the complexity.

Create an evidence-backed AI roadmap

A credible roadmap should explain not only what the organization intends to deploy, but why. For each prioritized opportunity, leaders should be able to articulate:

- The business outcome
- The process being improved
- The expected operational effect
- The context AI will require
- The level of human oversight
- The risks and controls
- The success measures
- The basis for scaling

This creates a more defensible conversation with business sponsors, finance, risk and the board. It also reduces the tendency to treat every technically possible use case as an investment priority.

How ARIS supports the decision

ARIS Process Lens can support evidence-led prioritization by analyzing governed process content for indicators that support or constrain automation and agentic potential. Its role in this narrative is not to select technology on behalf of the organization. It is to help turn process evidence into a more transparent and defensible backlog of opportunities.

The business value of better prioritization

Evidence-led prioritization concentrates funding on higher-value opportunities, reduces expenditure on use cases that should not progress and identifies process or context conditions that must be corrected before deployment.

It also creates a repeatable method for evaluating new opportunities as demand grows. Simulation provides the ability to quickly and easily understand the potential benefits and KPI impacts of changes to process execution.

Chapter takeaway:

The objective is not to reduce experimentation. It is to ensure experimentation produces evidence, and that evidence governs where the organization scales.

Chapter 4:

Equip AI to Operate Reliably

A capable model does not automatically become a reliable enterprise operator. Reliability depends on understanding the conditions under which AI is expected to act.

Many AI discussions focus on what agents can do. Far fewer focus on how agents should behave. That distinction becomes critical in enterprise environments.

The difference between capability and operational reliability

An AI agent may be capable of assessing a request, generating a recommendation or initiating an action. However, enterprise execution requires additional understanding:

- ❑ Is the agent authorised to make the decision?
- ❑ Which policy applies in this situation?
- ❑ Does the process differ by jurisdiction, product or customer type?
- ❑ Which upstream conditions must be satisfied?
- ❑ What downstream consequences could the action create?
- ❑ What level of confidence is acceptable?
- ❑ Which exceptions require human judgement?
- ❑ What evidence must be retained?
- ❑ When must the agent stop and escalate?

Without this context, the agent may produce a plausible action that is inconsistent with the organization's operating requirements. This is the difference between an AI system that can perform a task and one that can participate reliably in enterprise operations.

Governance must become operational

Many organizations have established high-level AI principles. These may cover fairness, transparency, security, human oversight and accountability. Such principles are essential, but they must be translated into the specific context of each process.

Operational governance defines:

- ☐ What the AI is permitted to do
- ☐ What the AI is not permitted to do
- ☐ Which data and knowledge sources are approved
- ☐ What confidence level is required
- ☐ Which decisions require human approval
- ☐ Which exceptions trigger escalation
- ☐ What information must be recorded
- ☐ How accountability is maintained
- ☐ How performance and compliance will be monitored

These requirements should not exist separately from the process—they are part of how the process is designed to operate. Embedding them into Process Context helps turn governance from a set of general statements into actionable operating boundaries.

Make approved context available when it is needed

Creating Process Context is only part of the answer. AI must also be able to access the relevant context at the point of action.

That context may include:

- ☐ Approved process definitions
- ☐ Standard operating procedures
- ☐ Decision rules
- ☐ Policies and controls
- ☐ Role definitions
- ☐ Organizational terminology
- ☐ Required hand-offs
- ☐ Exception paths
- ☐ Escalation conditions

Providing context at runtime helps connected AI understand how a requested action fits into the wider operation. It can also reduce the repeated effort involved in reconstructing context for every interaction or deployment.

The objective is not to give every agent unrestricted access to all enterprise knowledge. It is to provide controlled access to the approved context relevant to the task and process.

How ARIS supports runtime context

ARIS MCP provides a controlled mechanism for making governed process models, rules, policies and related context available to connected AI when needed. MCP does not execute the process or replace the agent platform. It provides a standardised connection to the approved operational understanding held in ARIS.

Design for human and AI collaboration

Reliable AI does not mean removing people from every decision. The strongest operating models deliberately define how responsibility is shared. AI may prepare a recommendation, gather and structure evidence, complete a low-risk activity, monitor process conditions, detect an exception or initiate a standard response.

A person may approve a high-impact decision, resolve an ambiguous exception, exercise regulated judgement, override an automated action, update the operating rule or accept accountability for the outcome.

These roles should be designed, not discovered after deployment. Process Context provides the structure needed to define where AI can act, where people remain accountable and how work moves between them.

Reliability as a prerequisite for scale

As the number of agents grows, informal oversight becomes less viable. Enterprise scale requires repeatable operating boundaries. A shared Process Context foundation supports this by helping organizations provide consistent rules, responsibilities and escalation logic across deployments.

Chapter takeaway:

Reliability is not a constraint on AI value. It is what makes sustained value possible.

CASE STUDY

Giving AI the Context to Operate Reliably

For AI agents to act autonomously inside complex enterprises, they need more than access to data. They need to understand how the business is supposed to work—including the processes, rules and boundaries within which they can act.

Global aerospace, defense and security company Leonardo is building that foundation with ARIS. Across its five divisions, Leonardo has created a governed Digital Twin containing more than 5,000 process models, with its 62,000+ employees able to access the process landscape.

Leonardo is also using Process Mining to connect this designed view with how work actually happens, identifying deviations, bottlenecks and opportunities to improve complex engineering and product-development processes.

Now that foundation is being extended to AI. Leonardo plans to use its Digital Twin to provide agents with the Process Context, rules and operating boundaries they need to act autonomously but reliably—grounding AI in an approved understanding of the business.

The lesson is clear: before AI can transform how work gets done, it needs a trusted understanding of the business it is being asked to transform.



“If we can influence the agents’ design by giving them the necessary process context and intelligence, that is a very important objective.”

Massimo Scalvenzi
Head of Quality Solutions,
Leonardo

Chapter 5:

Make AI Performance Visible

AI programmes cannot be managed effectively through deployment metrics alone. Leaders need to see what changes inside the operation after AI is introduced.

The number of users, prompts, agents or completed tasks may demonstrate adoption, but they do not demonstrate business value. Instead, leaders must look beyond mere data to understand how processes and workflows have changed.

Move beyond output measures

An AI output can be evaluated for speed, accuracy or completion. Those measures are useful but incomplete. Consider an agent that resolves a customer request faster. If the resolution creates additional rework, increases repeat contact or introduces a compliance exception, the initial output metric presents only part of the story.

The more important questions are:

- ☐ Did the end-to-end process perform better?
- ☐ Did the customer receive a better outcome?
- ☐ Was total handling effort reduced?
- ☐ Did cycle time improve?
- ☐ Were fewer exceptions created?
- ☐ Did the change affect another function or process stage?
- ☐ Did the organization reduce cost or release capacity?
- ☐ Did risk increase or decrease?

These questions require a process-level view of performance.

Connect AI activity to execution

To understand AI's contribution, organizations need to connect agent actions with the operational events and business objects surrounding them.


That means seeing AI activity in relation to:

- Cases, orders, requests or transactions
- Customers, suppliers, products or assets
- Upstream events
- Downstream outcomes
- Human interventions
- Process variations
- Exceptions and rework
- Control results
- Time, cost and quality measures

This provides a richer picture than an isolated AI dashboard. It shows not only what the agent did, but what happened because of it.

Establish an outcome hierarchy

AI measures should connect technical activity to business performance. A practical hierarchy may include four levels.



Adoption may be necessary, but it is not the final outcome. High usage should never be confused with high value.

Understand under-performance, not only success

Visibility is also essential when an initiative does not deliver the expected return. Possible causes include the wrong process being selected, insufficient context, inconsistent process inputs, high exception rates, problematic human hand-offs or a weak operating design.

Without execution visibility, teams may respond by adjusting the model when the real issue exists elsewhere in the process. Connecting AI performance to operational reality helps leadership direct corrective action towards the actual constraint.

How ARIS supports execution insight

ARIS Process Mining (OCPM) can connect operational events and business objects to reconstruct how end-to-end work actually runs. Where AI activity can be linked to that execution, leaders can assess its contribution within the wider process, identify exceptions and understand downstream effects that output dashboards may not reveal.

Create accountability for AI outcomes

As AI becomes part of business execution, ownership must extend beyond the technology team. A sustainable model requires clarity about who owns the business outcome, who owns the process, who owns the AI capability, who monitors risk and control performance, who approves operating changes and who responds when performance declines.

Performance visibility provides the shared evidence these stakeholders need. It creates a common basis for decisions about investment, intervention, redesign and scale.

Chapter takeaway:

Do not ask only whether the agent completed the task—ask whether the process performed better.

Chapter 6:

Creating a Continuous AI Improvement Model

Enterprise AI is not a sequence of deployments. It is an evolving operating model that must learn as the business changes.

Models change. Processes change. Regulations change. Customer expectations change. The business context surrounding an AI action can change even when the underlying technology does not.

Sustained value therefore depends on the organization's ability to learn and adapt.

Context cannot remain static

The context used to design or configure an AI initiative represents the organization at a particular moment. Over time, policies are updated, responsibilities move, systems are replaced, new products are introduced, process variants emerge, regulations change, performance priorities shift and new exceptions appear.

If AI continues to operate using outdated assumptions, the gap between intended and actual execution grows. Maintaining Process Context is therefore part of AI operations, not simply part of implementation.

Close the loop between design and execution

A continuous improvement model connects five activities.

Design: define how the process should operate, including roles, rules, controls, decisions and expected outcomes.

Deploy: introduce AI where evidence indicates it can improve performance, with clear operating boundaries and measures.


Observe: monitor execution across people, systems and AI to understand what is happening in the end-to-end process.

Improve: use evidence to address bottlenecks, exceptions, governance gaps and opportunities for further automation.

Update: reflect approved changes in the operational foundation so future decisions and AI activity use current context.

This creates a learning system. Execution improves the context. Better context improves future execution.

Make improvement portfolio-wide



The greatest value emerges when learning is not confined to one deployment. Evidence from one process can strengthen future opportunity assessment, agent design standards, governance patterns, reusable controls, context requirements, human hand-off models, performance measures and technology architecture.

This makes each initiative more valuable than its immediate result. The organization is not only deploying AI. It is improving its ability to deploy AI.

Create a shared management discipline

Continuous improvement requires coordinated ownership. CIOs and AI leaders provide technology strategy, architecture and platform governance. Line-of-business leaders own business outcomes and determine where operational change creates value. Process leaders connect local change to end-to-end performance. Risk and compliance teams define operating requirements and monitor control effectiveness. Data and technology teams ensure the necessary information and services are available.

A shared Process Context foundation gives these stakeholders a common reference point. Different teams may have different responsibilities, but they are working from the same understanding of the operation.

How ARIS helps maintain and improve the foundation

ARIS AI Companion can help teams accelerate the creation and exploration of structured process knowledge and support analysis across design and mining environments. Combined with the Digital Twin and execution insight, this can reduce the effort required to keep Process Context useful as the organization changes.

The objective is not to automate governance or remove expert judgement. It is to make the work of creating, using and maintaining trusted operational understanding more efficient and repeatable.

Process Overlay makes it simple to compare approved, governed ways of working with operational reality. It highlights opportunities to improve the efficiency and effectiveness of processes and their associated controls. It also exposes risks arising from non-conformance and agent drift.

The compounding effect

A fragmented AI programme starts again with every new use case. A mature AI operating model reuses what the

organization has already learned. It reuses process knowledge, decision rules, governance patterns, system relationships, performance evidence, exception insights and operating controls.

This changes the economics of scale. The cost and effort of every new initiative no longer need to rise at the same rate as adoption. The foundation becomes more valuable as it is reused and improved.

Chapter takeaway:

That's how AI value begins to compound: not by deploying more in isolation, but by improving the organization's ability to choose, govern, measure and refine every deployment.

CASE STUDY

Putting Process Context to Work with AI Agents



L&G has built a governed repository of more than 2,000 standardized business processes in ARIS, creating a trusted source of knowledge covering applications, risks, controls, organizational responsibilities and dependencies.

Crucially, this process knowledge is no longer simply documentation—double-digit numbers of AI agents are already using ARIS process models as part of their operational context.

Because those processes are reviewed, approved and maintained by the business, agents can work from trusted Process Context rather than disconnected documents or outdated information. This is already helping L&G accelerate business analysis delivery across transformation projects by 20%.

L&G is also using AI to accelerate process modeling, generating draft models that subject-matter experts can review and approve—creating a powerful cycle in which AI consumes trusted process knowledge while also helping the organization continuously develop and improve it.

The result is a foundation that can evolve alongside the business: governed Process Context powering multiple agents today, while creating the platform for continuous AI-driven improvement tomorrow.

Conclusion:

Realizing Enterprise Value with AI

The organizations that scale AI successfully establish an operating foundation and management discipline that turn capability into sustained business performance.

Scaling AI is not primarily about adding more models, platforms, tools or agents. It requires an operating foundation and management discipline that help leaders make four decisions well.

Focus investment

Direct funding and organizational attention towards processes where AI can create meaningful operational and financial value.

Equip AI to operate reliably

Provide the approved context, operating rules, controls and escalation conditions AI needs to participate safely and effectively in enterprise operations.

Make performance visible

Connect AI activity to end-to-end process and business performance so leaders can distinguish adoption from value.

Improve continuously

Use execution evidence to refine processes, update context, strengthen governance and improve future AI investment decisions.

Process Context connects these priorities. It gives the organization a shared understanding of how work is designed, how work is executed and how AI should participate. The Digital Twin provides the foundation for that understanding. It brings together business design and operational reality so that improvement and AI scale can be managed as connected outcomes.

The implications for enterprise leaders

- ❑ AI teams do not need to reconstruct the enterprise for every use case.
- ❑ Business leaders gain a clearer connection between AI investment and operational performance.
- ❑ Governance becomes part of execution rather than an activity applied after deployment.
- ❑ Performance can be evaluated through business outcomes rather than technical outputs alone.
- ❑ Operational evidence can continuously improve both the process and the AI operating within it.

This is the difference between an organization that deploys AI and an organization that develops the ability to scale AI. The first creates a growing collection of initiatives. The second creates a repeatable system for converting AI investment into enterprise value.

That system becomes increasingly important as AI capability becomes more widely available. Models, platforms and agents will continue to advance. Access to AI technology will become less differentiating over time.

The greater source of advantage will be the organization's ability to apply AI to the right opportunities, within the right operating context, with clear visibility into performance and a disciplined method for continuous improvement.

Final perspective

That is the foundation for scalable enterprise AI. And it is how AI moves from technological potential to sustained business performance.

Where ARIS fits

ARIS provides the Process Context Platform for Enterprise AI, bringing together the designed and observed views of how the enterprise operates in a connected Digital Twin.

The platform can help organizations:

- Build and maintain trusted Process Context
- Identify and prioritize high-value opportunities
- Make governed process context available to connected AI
- Understand operational execution across end-to-end processes
- Connect insight to continuous improvement

**Explore how ARIS can help
your organization establish
the Process Context required
to scale AI with greater focus,
reliability and visibility.**

[Book a call with an expert](#)



ARIS is the process context foundation platform for enterprise AI deployment. Combining process mining, modelling, and analysis in a single unified platform, ARIS helps leading global organizations move from AI experimentation to scalable execution – driving efficiency, reducing risk, and delivering measurable business outcomes.

© 2026 SAG ARIS GmbH. All rights reserved.

The name ARIS and SAG ARIS GmbH product names are either trademarks or registered trademarks of SAG ARIS GmbH and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

ebook-aris-why-enterprise-ai-stalls-en