



The pragmatist's guide to enterprise AI:

Why the Activate, Build, and Connect approach outperforms Overlay approaches in today's complex IT landscapes



Purpose of the document

This position paper is designed for CEOs, CFOs, COOs, and other senior business leaders navigating the strategic implications of AI in global enterprises.

Throughout this paper, we refer to “**Overlay**” approaches as third-party AI platforms that sit above existing business applications, extracting data to a central repository for AI orchestration. This contrasts with the “**Activate, Build, and Connect**” approach, where organizations start with AI capabilities embedded directly into business applications, progressively build a semantic data foundation, and connect specialized agents across systems.

While it’s widely acknowledged that AI-enabled orchestration through agentic scenarios drives incremental gains in efficiency—translating them into tangible outcomes such as improved working capital, revenue growth, and margin expansion—this paper presumes that senior leaders are already aligned on these benefits. Consequently, this paper focuses instead on the strategic and operational pathways that determine whether you can realize such benefits at scale and sustain them over time.

The purpose of this document is to:

1. **Clarify the strategic choice** between the Activate, Build, and Connect approach from software vendors versus third-party Overlay approaches of AI platforms in today’s complex IT environments.
2. **Demystify the hidden technical and implementation realities** that determine AI success but which are rarely discussed in executive-level conversations. These include both the architectural challenges of preserving business context (see p23 for key concept definitions) and the organizational factors that enable high-performing organizations to achieve significantly better results through strong AI data readiness.
3. **Provide a progressive, educational narrative** that defines key concepts and builds toward a research-based conclusion about the optimal path to enterprise-wide AI value.
4. **Present the Activate, Build, and Connect roadmap** for realizing the vision of end-to-end process intelligence while addressing the six major dimensions that separate successful implementations from failed pilots.

While this document reflects SAP’s perspective on enterprise AI strategy, it is grounded in independent research and addresses the reality of today’s hybrid application landscapes. Our aim is to help senior leaders make informed decisions that deliver measurable business value while avoiding the costly pitfalls that have characterized many ambitious AI transformation efforts.



Agenda

01

The Challenge of Enterprise AI
Executive Summary

02

The Evolution of Enterprise IT
From Monolith to Mosaic

03

Enterprise AI
Understanding the Two Approaches

04

The Hidden Reality
Why “Overlay” Approaches Fail

05

The Progressive Path
“Activate, Build, and Connect” approach

06

Making the Strategic Choice
Decision Framework

07

The Pragmatic Path Forward
Conclusion

1. The challenge of enterprise AI

Executive summary

1.1. The vision: End-to-end process intelligence

Today's executives envision a compelling future: where AI agents orchestrate seamless end-to-end business processes across their increasingly complex application landscapes. This vision promises to transcend the fragmentation that has evolved as corporate IT environments have shifted from monolithic ERPs to specialized best in class solutions. It's an attractive destination that aligns with the strategic imperative to drive enterprise-wide efficiency and innovation.

1.2. The challenge: Best in class vs. end-to-end intelligence

A fundamental challenge has emerged in enterprise AI implementation. The very evolution that has delivered tremendous business user satisfaction—moving from monolithic ERPs to best in class applications—has created significant barriers to achieving end-to-end AI orchestration. Ironically, if companies had maintained wall-to-wall ERP systems, implementing end-to-end AI would be simpler. Instead, today's fragmented application landscapes require a more sophisticated approach.

1.3. The implementation crisis

Many organizations recognize the transformative potential of AI but face a fundamental challenge in realizing that value. Recent research across multiple studies reveals that while AI adoption is widespread, only a minority of organizations achieve measurable business impact. Two realities explain why data readiness and implementation approach have become the defining factors:

First reality: The business context challenge. When data is extracted from enterprise systems and replicated to generic data lakes, it loses the rich semantic meaning that makes it valuable. Research confirms that the majority of executives cite this "loss of business context" as a major barrier to AI success, a finding consistent across multiple independent studies. Rebuilding this context is prohibitively expensive and time-consuming and not successful in most cases, creating an insurmountable barrier for many AI initiatives.

Second Reality: The pilot-to-production chasm. Research from Massachusetts Institute of Technology (MIT)¹ reveals a consistent pattern: while 60% of organizations evaluate enterprise AI tools and 20% reach pilot stage, **only 5% achieve production deployment with measurable profit-and-loss (P&L) impact.** This gap exists not because of technical limitations, but because organizations underestimate six critical dimensions of production readiness: learning capability, workflow integration, change management, data governance, maintenance burden, and measurable business outcomes.

Together, these realities explain the data readiness confidence gap identified by Oxford Economics research²: **55% of executives lack confidence in their ability to integrate data across business functions, 75% face incomplete or inconsistent data, and 68% struggle with siloed data.** Yet the same research shows that organizations with high AI data readiness achieve significantly better business value—demonstrating that the path to success lies not in technological limitations, but in how organizations approach data foundation and implementation strategy. The Activate, Build, and Connect approach directly addresses these readiness challenges by preserving business context and building progressively toward end-to-end intelligence.

1.4. The path forward: Activate, Build, and Connect

The research suggests that organizations that successfully achieve end-to-end process intelligence follow a different path than those pursuing centralized "single brain" architectures. Rather than attempting to build a monolithic AI layer that sits above all systems, successful organizations:

- 1. Activate domain-specific intelligence:** Initiate embedded AI capabilities within core business applications where business context is the richest and most valuable.
- 2. Build cross-system intelligence:** Establish a semantic data fabric that preserves business context across systems while enabling cross-application intelligence.
- 3. Connect through agent-to-agent communication:** Create a network of specialized AI agents that communicate across application boundaries rather than a single orchestration layer.

If you follow all three phases, it should allow you to address the six dimensions of **production readiness**: you will ensure AI systems can learn and adapt, integrate deeply into workflows, secure executive sponsorship and user adoption, maintain clear data governance, minimize maintenance burden, and demonstrate measurable business outcomes.

This Activate, Build, and Connect approach delivers the end-to-end process intelligence executives seek through a pragmatic path that generates value at each step. By addressing both the architectural challenges of business context preservation and the organizational challenges of scaling from pilot to production, this approach enables organizations to join the high-performing minority that successfully transforms AI investments into sustained business value.

2. The evolution of enterprise IT

From monolith to mosaic

2.1. The historical Shift: Wall-to-wall ERP to best in class

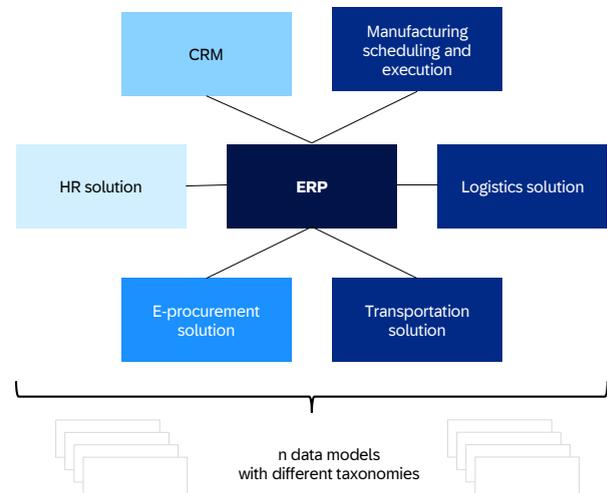
Enterprise IT landscapes have undergone a profound transformation over the past two decades. What was once a single, integrated ERP system covering nearly all business processes has evolved into a complex ecosystem of specialized applications—each optimized for specific business functions:

- **Past (1990s-2000s):** End-to-end processes executed almost exclusively within a single ERP system with one unified data model
- **Present (2010s-2020s):** Virtually no end-to-end process executes within a single application; instead, processes flow through multiple specialized systems (such as ERP, customer relationship management (CRM), human capital management, supply chain management, and so on.)

Corporate IT landscapes have evolved from this:



To this:



This evolution has delivered tremendous business user satisfaction through specialized capabilities but has created new challenges for process integration and data coherence.

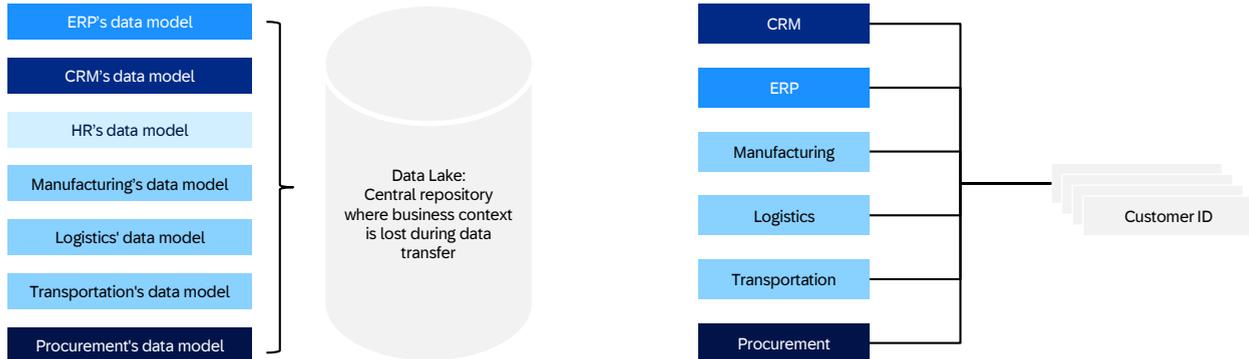
2.2. The data integration challenge

This fragmentation has fundamentally changed how enterprises manage their data.

- **Past:** Data was replicated from the ERP to a data warehouse primarily for performance reasons, and to build operational reporting, but the semantic relationships remained intact.
- **Present:** Data must be extracted from multiple systems with different data models, taxonomies, and business rules, then consolidated in data lakes.

Moving data from various business systems into a centralized repository strips away essential business context.

This forces costly IT efforts to restore business context and standardize definitions before data can deliver cross-functional insights.



This shift has created an often-overlooked challenge: the loss of business context when data is moved from its native environment to centralized repositories.

2.3. The AI implementation challenge

The fragmentation of enterprise applications creates both a challenge and an opportunity for AI implementation:

- **The vision:** Many organizations seek end-to-end process orchestration across all business applications.
- **The success pattern:** Research shows that organizations with high AI data readiness achieve significantly better business value
- **The challenge:** Many organizations struggle because contextually rich, semantically harmonized data remains difficult to create in fragmented landscapes.

3. Enterprise AI

Overlay vs. Activate, Build, and Connect

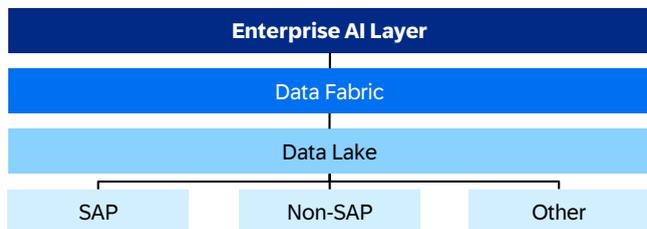
The fragmentation of enterprise applications has created an urgent need for cross-system AI orchestration. Yet the path forward is not obvious. Two fundamentally different approaches have emerged, each with distinct architectures, promises, and risk profiles. Understanding this choice is critical; it will determine whether your organization joins **the 5% that successfully scale AI or the 95% that struggle¹**. The question is not whether to pursue enterprise AI, but which architectural approach will actually deliver on the vision of end-to-end process intelligence.

3.1. Architectural approach

The Overlay approach

The first approach, often promoted by technology platform providers and consulting firms, proposes creating a centralized AI layer that sits above all enterprise applications.

AI Use Cases : 100% custom



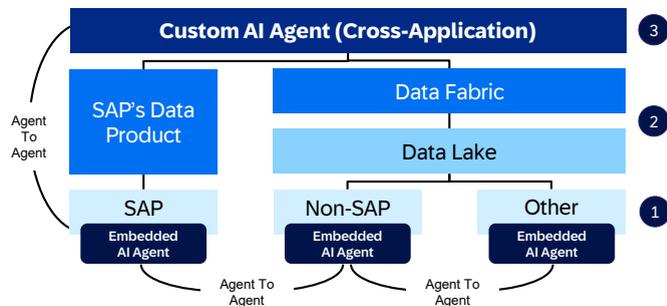
Definition: A third-party AI platform that connects to various business applications, extracts their data, and attempts to orchestrate processes across system boundaries.

Promise: A single brain that can understand and orchestrate end-to-end processes regardless of the underlying applications.

The Activate, Build, and Connect approach

The second approach, typically advocated by enterprise software vendors, proposes adopting a phased approach.

AI Use Cases : 80% embedded and 20% custom



Definition: First, utilize AI capabilities embedded within business applications (1, see above), then build a hybrid data platform (2) to enable cross-application AI agents (3) while leveraging agent-to-agent communication.

Promise: Domain-specific intelligence that delivers immediate value while building toward end-to-end orchestration.

3.2. The difference

Dimension	Overlay Approach <small>AI Use Cases: 100% custom</small>	Activate, Build, and Connect <small>AI Use Cases: 80% embedded and 20% custom</small>
Starting Point	Build central AI orchestration layer first	Activate embedded AI capabilities first
Business Context	Must rebuild in central repository	Preserved in native applications
Time to First Value	Months to years (after data integration)	Days to weeks (immediate activation)
Risk Profile	High (big-bang transformation)	Low (incremental implementation)
Maintenance Burden	All-in: Customer maintains the entire semantic layer	Shared: Vendor maintains embedded capabilities and data product while customer maintains custom semantic layer
Path to End-to-End Intelligence	Direct but rarely achieved	Progressive and proven
User Adoption	AI as a separate layer, requiring extra training.	AI embedded in workflows, enabling natural adoption.

The Overlay approach promises a shortcut to end-to-end orchestration. The Activate, Build, and Connect approach acknowledges that there are no shortcuts, only a pragmatic path that delivers value at each stage.

The choice between these approaches will determine your AI success.

4. The hidden reality

Why Overlay approaches fail

The Overlay approach's conceptual simplicity—where one AI layer orchestrates all systems—is appealing. However, two hidden realities consistently prevent this vision from becoming reality. These challenges are rarely discussed in executive presentations, yet they explain why the vast majority of Overlay implementations fail to deliver measurable business value.

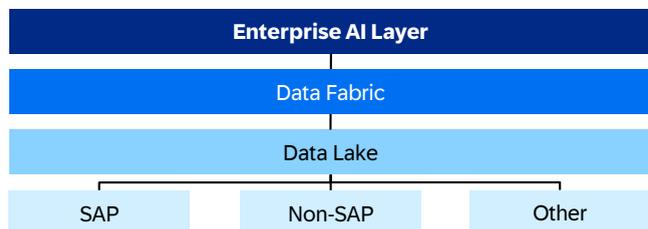
4.1. First hidden reality: The business context challenge

Overlay approaches face a fundamental architectural dilemma. To orchestrate processes across multiple systems, you must access data from those systems.

4.1.1. Two paths exist, both problematic

Path 1: Centralized data lake

Extract data from all business applications into a central repository for AI processing.



The challenge

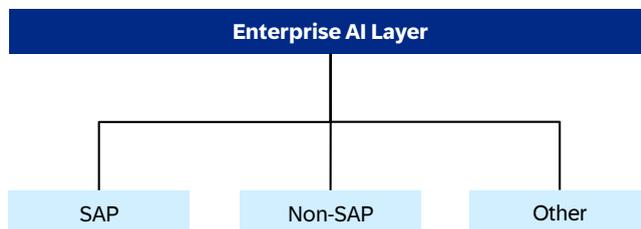
Semantic rebuilding: Before you can realize any AI value **at scale**, all business context must be rebuilt in the data fabric—a process that can take years and cost millions.

Maintenance burden: As source applications evolve, the rebuilt semantic layer must be continuously updated to maintain accuracy.

Time to value: Organizations must complete massive data integration projects before delivering any tangible AI benefits.

Path 2: Direct database connection

Connect the AI orchestration layer directly to the production application databases.



The challenge

Data orchestration: AI layers function as query interfaces, not integration platforms, which means they can't extract, transform, or synchronize data across SAP, non-SAP, and legacy systems—leaving critical gaps in data harmonization and consistency.

System saturation: Excessive data queries can overload transactional production systems, potentially causing downtime that could paralyze business operations.

Compliance issues: Direct access to the application data through the database layer often violates licensing terms and creates security vulnerabilities.

4.1.2. Understanding the business context challenge

When organizations choose Path 1 (centralized data lake), they face a fundamental problem: the critical business context disappears during data extraction.

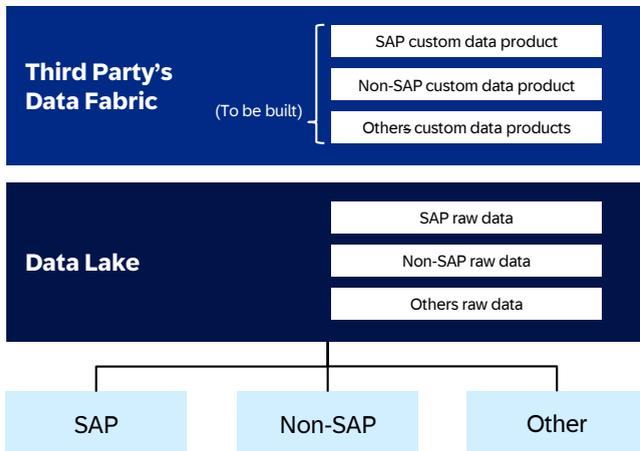
What's lost? Business attributes, hierarchies, currency conversions, units of measure, time-dependent relationships, and business rules that give data meaning.

The consequence: AI models trained on decontextualized data deliver generic answers rather than business-specific insights.

The restoration challenge: Organizations face two options for creating data products

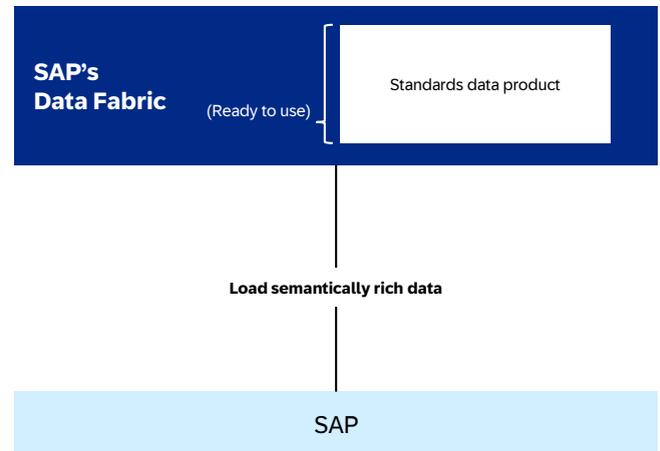
Build custom data products

This requires substantial data engineering resources to recreate the business context, plus ongoing maintenance as source applications evolve. This is the path most organizations must take when attempting Overlay approaches.



Utilize vendor-provided data products

Some software vendors (such as SAP and Salesforce) offer pre-built data products with business context already preserved, accelerating time to value for native applications while allowing selective custom development for strategic cross-system insights.



The reality

The choice between custom and vendor-provided data products significantly impacts time to value, maintenance burden, and total cost of ownership, factors we'll explore in the progressive approach presented in chapter 5.

4.1.3. Addressing the objection: “We’ve already invested millions”

Objection: “We’ve already invested millions in our data lake and data fabric. Are you saying we should abandon it?”

Response: Not necessarily. The answer depends on what you’ve actually built.

Scenario A—raw data Lake: Your data lake serves analytics and reporting well, where some context loss is acceptable. Continue using it for these purposes while pursuing embedded AI for process intelligence. These aren’t mutually exclusive.

Scenario B—Rebuilt semantic layer: If you’ve successfully rebuilt business context across all enterprise systems, the key questions are: How many full-time equivalents maintain the semantic layer? What percentage of business context is recreated? Can it support real-time orchestration? Can we sustain the semantic layer as applications evolve? If your answers show low maintenance, high completeness, and real-time capability, you may be positioned for AI orchestration—though fewer than 5% of organizations reach this state.

Scenario C—Hybrid approach: You likely have vendor data products (such as the SAP Business Data Cloud solution or Salesforce Data 360, formerly Data Cloud) plus custom data fabric for non-vendor systems.

The optimal strategy is to:

1. Use embedded AI for single-vendor domains.
2. Utilize vendor data products where they exist.
3. Use custom data fabric selectively for high-value cross-vendor scenarios.

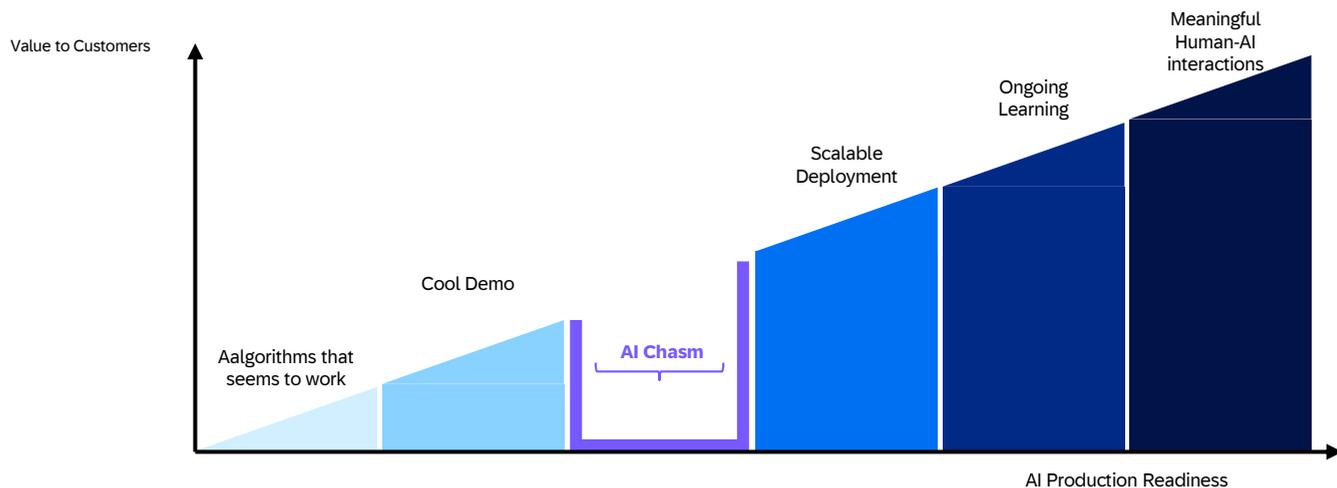
4.2. Second hidden reality: The pilot-to-production chasm

4.2.1. Most AI pilots never reach production scale

Even when organizations successfully address the architectural challenges above, a second barrier emerges: most AI pilots never reach production scale.

Research from MIT’s Project NANDA reveals a stark pattern: while 60% of organizations evaluate enterprise AI tools and 20% reach pilot stage, only 5% achieve production deployment with measurable P&L impact. This 75% drop-off occurs not because of technical limitations, but because organizations underestimate what it takes to scale from proof-of-concept to sustained business impact.

This implementation chasm affects both Overlay and Activate, Build, and Connect, but the barriers differ dramatically. Six major dimensions of production readiness separate successful implementations from failures.



4.2.2. The six dimensions of production readiness

1. Learning and adaptation capability

The most common failure mode is brittleness, where systems work in controlled tests but break in edge cases. According to MIT's 2025 research¹, two thirds of executives cite “the ability to improve over time” as their top vendor selection criterion, yet most enterprise AI tools remain static, unable to adapt as business processes evolve.

Embedded AI advantage: Vendor-maintained capabilities evolve continuously as applications evolve, eliminating the burden of retraining custom models when business processes change.

2. Workflow integration depth

Pilots succeed in isolation but fail when integrated into complex, multisystem workflows. The difference between 20% pilot adoption and 5% production deployment is typically integration complexity.

Embedded AI advantage: AI capabilities built directly into business applications are already integrated by design. For example, an accounts receivable agent in SAP S/4HANA doesn't need separate integration with payment processing. Overlay approaches require building and maintaining these integrations manually.

3. Change management and user adoption

MIT researchers identify “challenging change management” and “unwillingness to adopt new tools”¹ as top barriers to scaling AI. Users who enthusiastically adopt ChatGPT for personal use often resist enterprise AI tools that don't match their expectations for responsiveness and quality.

Embedded AI advantage: Activating AI within familiar applications dramatically reduces change management friction, because users gain new capabilities within systems they already use daily rather than learning entirely new interfaces.

4. Data governance and security boundaries

As AI moves to production, governance concerns intensify. MIT's survey shows 61% of executives cite "clear data boundaries" as critical, particularly in regulated industries where client data cannot be mixed with other organizations' training data.

Embedded AI advantage: AI capabilities within business applications inherit existing security models, compliance certifications, and governance frameworks. Overlay approaches require building parallel governance structures, often delaying production deployment by 6-12 months.

5. Maintenance burden and total cost of ownership

Pilots often ignore ongoing maintenance costs. This is particularly acute for Overlay approaches requiring continuous semantic layer maintenance as source applications evolve.

Embedded AI advantage: Vendor-maintained embedded AI shifts this burden from customer to vendor. When SAP releases a new S/4HANA version, embedded AI updates automatically. With Overlay approaches, each source application upgrade can break semantic mappings, requiring manual remediation.

6. Measurable business outcomes

The final barrier is demonstrating measurable P&L impact, not just productivity improvements. Organizations that successfully scale AI move beyond benefits like faster email responses to quantifiable business metrics such as reduced days sales outstanding (DSO) and lower inventory carrying costs.

Embedded AI advantage: Because embedded AI operates within transactional systems, business outcomes are directly measurable and tie clearly to financial statements.

4.2.3. The implementation advantage

Building on MIT's research on implementation success rates, an implicit pattern emerges: organizations build external partnerships with vendors providing learning-capable, embedded AI—which represents a good starting point in their enterprise AI journey.

This implementation advantage stems from three factors:

1. Domain expertise from thousands of implementations across industries.
2. Continuous evolution that improves capabilities across all customers simultaneously.
3. Proven implementation playbooks that reduce costly trial and error.

These two hidden realities, the business context challenge and the pilot-to-production chasm, explain why Overlay approaches consistently fail despite their conceptual appeal.

5. The progressive path

Activate, Build, and Connect

The previous section revealed why Overlay approaches fail: they don't allow organizations to preserve business context, they make it difficult for organizations to rebuild and maintain it, they underestimate the six dimensions of production readiness. This section presents the alternative, a progressive path that directly addresses both hidden realities while delivering measurable value at each stage. Rather than attempting to build end-to-end AI orchestration in a single leap, successful organizations follow three sequential principles that build upon each other.

5.1. The three sequential principles

5.1.1. Step 1: Start where the context is by activating embedded AI capabilities

The fastest path to AI value begins by activating capabilities already built into your core business applications:

AI Use Cases : 100% embedded



Definition: Embedded AI refers to intelligence capabilities built directly into business applications by their vendors to utilize the full business context of the application, including its data model, business rules, and process flows.

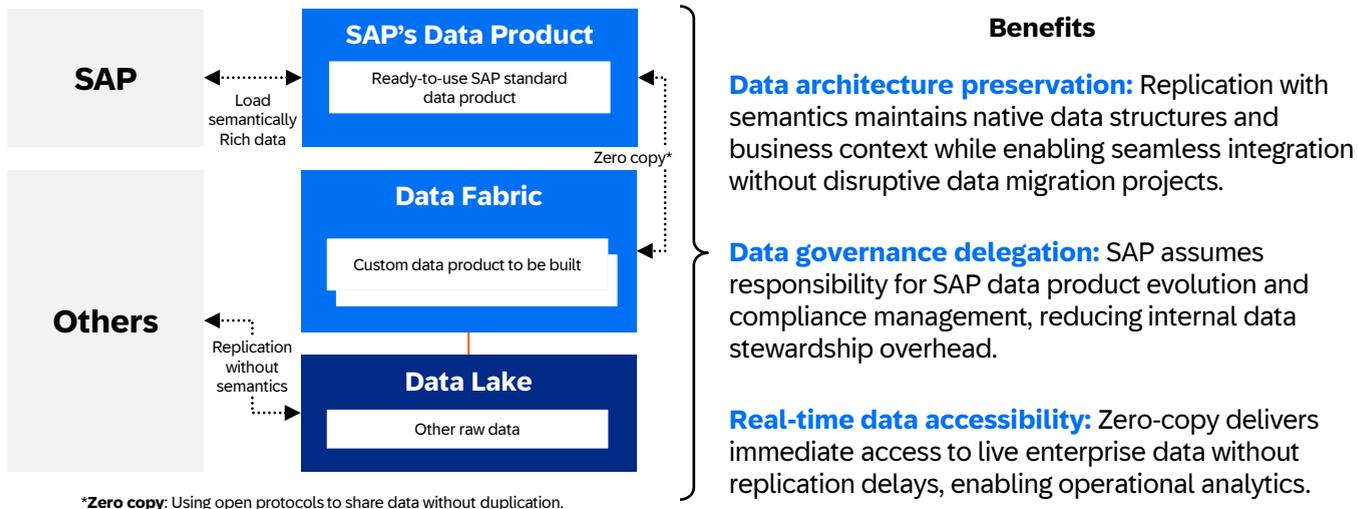
Example: An accounts receivable dispute resolution agent in SAP S/4HANA understands SAP's universal journal (unified financial ledger), as well as payment terms, customer history, and dunning procedures without requiring any context rebuilding.

Three immediate benefits:

- 1. Rapid time to value:** You can activate capabilities in weeks rather than the months or years required for data lake projects.
- 2. Built-in governance:** Security, compliance, and data privacy are handled within the application's existing framework, so no parallel governance structures are required.
- 3. Continuous evolution:** The vendor maintains and enhances these capabilities as the application evolves, eliminating customer maintenance burden.

5.1.2. Step 2: Build the data foundation

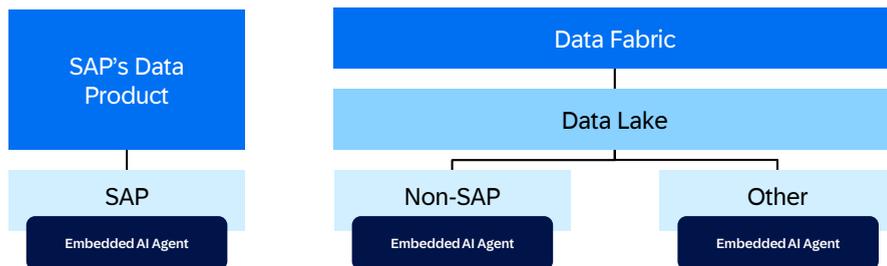
Virtually all global enterprises operate mixed vendor environments. A hybrid federation approach offers the most pragmatic path forward.



The important distinction: A data fabric provides the infrastructure for federating distributed data, while data products are the curated datasets that sit on top of that infrastructure.

SAP delivers both:

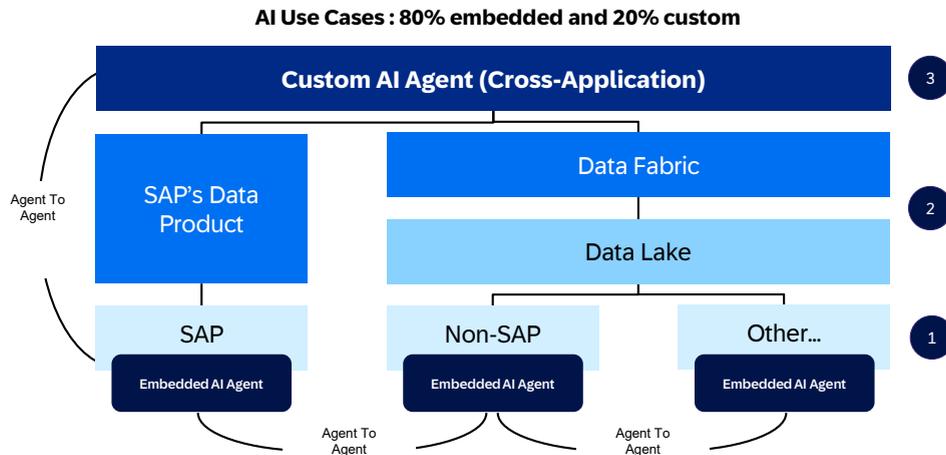
- The data fabric
- Ready-to-use data products (pre-built semantic layers with business context preserved)



Third-party approaches typically provide only the data fabric, requiring you to build the data products yourself, a process that can take years and cost millions of dollars.

5.1.3. Step 3: Evolve toward agent-to-agent communication

As embedded AI capabilities mature and the data fabric establishes cross-system connectivity, organizations naturally progress to enabling domain-specific agents to communicate directly.



Definition: Agent-to-agent communication allows AI capabilities from different applications to collaborate on end-to-end processes while maintaining their specialized domain expertise.

Examples: A sales order agent in a CRM system communicating with an inventory allocation agent in an ERP or a logistics optimization agent in a transportation management solution orchestrating the order-to-cash processes.

The federated advantage: This approach acknowledges the reality of heterogeneous application landscapes. Rather than forcing all intelligence into a single centralized brain, it creates a network of specialized agents, each maintaining deep domain expertise while contributing to cross-functional processes.

Why this works: Each agent operates where business context is richest (within its native application) and communicates via standardized protocols like agent-to-agent or model context protocol. For cross-application use cases, custom agents utilize the data fabric while maintaining interoperability with prebuilt vendor agents. The AI agent hub in SAP LeanIX Application Portfolio Management provides centralized visibility and policy-based control across all AI agents—whether that’s the Joule copilot, custom-built agents or third-party agents.

The progressive advantage

This phased approach delivers measurable value at each stage while progressively building the capabilities required for end-to-end process intelligence. Organizations that attempt to cross the chasm in a single leap, implementing “Overlay” AI for end-to-end orchestration from day one, consistently fail because they underestimate the six dimensions of production readiness.

5.2. The implementation timeline: Crossing the chasm

Organizations that successfully scale AI from pilot to production follow a disciplined, phased approach that addresses the six dimensions of production readiness at each stage. Here are the four timed phases for achieving success:

Months 1-3: Activate with clear metrics

- Select embedded AI capabilities with clear, measurable business outcome metrics (not just productivity improvements).
- Establish baseline measurements tied to P&L impact (such as reductions in DSO and inventory carrying cost).
- Involve frontline users in the pilot design to ensure workflow integration and reduce change management friction.

Months 4-9: Expand with learning loops

- Implement feedback mechanisms that enable AI to make improvements based on user corrections—to address the learning capability dimension.
- Monitor adoption metrics and user satisfaction to validate change management success.
- Demonstrate measurable business impact to secure executive sponsorship for a broader rollout.
- Begin building the data fabric foundation in parallel.

Months 10-18: Scale with governance

- Extend successful use cases to additional business units, taking advantage of proven implementation playbooks.
- Establish data governance frameworks that define clear boundaries and compliance requirements.
- Build toward cross-application scenarios as the data fabric matures.
- Measure and communicate cumulative P&L impact to maintain executive support.

Months 19+: Connect for end-to-end intelligence

- Utilize the data fabric to enable cross-application AI agents for high-value end-to-end processes.
- Implement agent-to-agent communication protocols for process orchestration.
- Continuously optimize processes based on measurable business outcomes.
- Expand to additional cross-functional scenarios as capabilities prove themselves.

6. Making the strategic choice

Decision framework

6.1. Five questions to guide your decision

When evaluating your enterprise AI strategy, the choice between Overlay and Activate, Build, and Connect approaches ultimately comes down to five critical questions:

Question 1

How quickly do we need to demonstrate measurable business value?

- **If you need scalable results within 90 days**, choose the Activate, Build, and Connect approach as it delivers immediate value through embedded AI capabilities that you can activate in weeks.
- **If you want to go beyond quick and dirty prototypes and you can wait 2-3 years for data integration** select the Overlay approach because you need to complete massive data lake projects before you can deliver any tangible AI benefits.

Question 2

What is our risk tolerance for AI investments?

- **If you need progressive de-risking**, opt for the Activate, Build, and Connect approach because it delivers value at each stage, allowing you to validate ROI before expanding investment.
- **If you can accept 95% failure probability**, pick the Overlay approach as it requires big-bang transformation with high upfront investment and years before seeing returns.

Question 3

Who will maintain the semantic layer as our applications evolve?

- **If you want vendor-maintained capabilities**, choose the Activate, Build, and Connect approach because it shifts the maintenance burden to vendors, which update embedded AI as applications evolve.
- **If you have data engineering teams available**, select the Overlay approach, which requires continuous updating of custom semantic layers as source applications change.

Question 4

How critical is preserving native business context for your AI use cases?

- **If business context is critical**, opt for the Activate, Build, and Connect approach as it preserves context by operating where it's richest—within native applications.
- **If generic insights are sufficient**, pick the Overlay approach because it accepts context loss during data extraction and attempts to rebuild it in centralized repositories.

Question 5

What is your path to end-to-end process intelligence?

- **If you want a progressive, proven path**, the Activate, Build, and Connect approach is for you because it builds toward end-to-end orchestration through phased implementation with value at each stage.
- **If you believe a direct path is achievable**, choose the Overlay approach because, even though it promises immediate end-to-end orchestration, it rarely delivers due to data complexity.

6.2. The reality

For most organizations, the answers to these questions point overwhelmingly toward the Activate, Build, and Connect approach. The Overlay approach may be appropriate only for organizations with:

- Exceptional data engineering capabilities.
- A tolerance for high risk, multi-million-dollar investments.
- The ability to cope with multi-year timelines before requiring an ROI.
- A willingness to maintain complex semantic layers indefinitely.

This combination is rarely found in practice, which explains why 95% of Overlay implementations fail to generate measurable financial returns.

7. The pragmatic path forward

Conclusion

7.1. Resolving the challenge

This paper began with a challenge: the very evolution that delivered tremendous user satisfaction, moving from monolithic ERPs to best in class-applications, has created significant barriers to achieving end-to-end AI orchestration.

The Overlay approach promises to resolve this challenge through a centralized AI layer that sits above all systems. It's a conceptually appealing vision that aligns with the executive's desire for unified process orchestration. However, as we've demonstrated, this approach forces organizations into high-risk, low-value data duplication projects that consistently fail due to two hidden realities: the impossibility of preserving business context in centralized repositories, and the underestimation of what it takes to cross the pilot-to-production chasm.

The Activate, Build, and Connect approach resolves the challenge differently, not by imposing a monolithic AI layer above fragmented systems, but by progressively building intelligence that respects the reality of heterogeneous application landscapes while delivering measurable value at each stage.

7.2. The research-backed path

This isn't merely SAP's perspective—it's a strategy that aligns with key findings from leading industry analysts and research institutions. The evidence on the challenges of AI transformation is clear and consistent. For instance, a recent September 2025 BCG study of over 1,250 firms highlights the difficulty of achieving returns on AI investment:

- 60% of companies are not achieving material value at all, reporting minimal gains despite substantial investment.¹
- Only 5% are achieving AI value at scale, demonstrating how tough the full transformation is.¹

These outcomes often stem from underestimating the complexity of preserving business context and a failure to address the critical dimensions of production readiness. Organizations that succeed follow a different path—the phased, progressive “Activate, Build, and Connect” approach outlined in this paper.

These failures share common characteristics: they underestimate the complexity of preserving business context, and they fail to address the six dimensions of production readiness that separate successful implementations from failed pilots.

Organizations that succeed follow a different path—the phased, progressive approach (Activate, Build, and Connect) outlined in this paper.

7.3. The executive imperative

As a senior business leader, your role is not to choose between technical implementations but to ensure your organization follows a path that:

1. Delivers measurable business value at each stage rather than requiring years of investment before seeing returns.
2. Minimizes risk and maximizes return on investment through incremental implementation that proves itself at each phase.
3. Builds toward your strategic vision in a sustainable way that doesn't compromise long-term goals for short-term wins.

The choice between Overlay and Activate, Build, and Connect is ultimately a choice between two fundamentally different risk profiles:



7.4. The path forward

The fragmentation of enterprise IT is not going away. Best in class applications will continue to proliferate as organizations seek specialized capabilities that deliver user satisfaction. The question is not whether to pursue end-to-end AI orchestration, but how to achieve it in a way that respects this reality.

The Overlay approach offers a shortcut that doesn't exist. The Activate, Build, and Connect approach offers a pragmatic path that does.

The challenge is resolved not by imposing centralized control over fragmented systems, but by building federated intelligence that takes advantage of the strengths of specialized applications while enabling them to work together seamlessly.

Organizations that recognize this reality, and act on it, will join the 5% that successfully transform AI investments into sustained, measurable business value. Those that chase the false promise of Overlay approaches will join the 95% that struggle.

The choice is yours. The evidence is clear. The path forward is proven.

Definitions

Key concepts defined

To understand the implications of this position paper, we must define several key concepts.

- **Business context or semantics (in data context):** The rich semantic relationships between data elements that give them meaning (for example, a sales order's connection to customers, products, pricing agreements, and fulfillment processes).
- **Data replication:** The process of copying data from source systems to a central repository, typically resulting in the loss of business context.
- **Data federation:** A technique that allows data to remain in source systems while making it accessible for cross-system analysis without physical movement.
- **Zero copy:** An open protocol for securely sharing data across organizations and platforms without duplicating it.
- **Data product:** A curated, contextualized dataset designed for specific business purposes, complete with metadata and semantic relationships.
- **Data fabric:** An architecture that facilitates the discovery, connection, and governance of distributed data across an enterprise.
- **Data lake:** A centralized repository that stores structured and unstructured data from multiple sources in its raw format, typically without the business context and relationships that existed in source systems. This differs from data warehouses, which maintain structured schemas.
- **AI Agent:** An autonomous software entity that can perceive its environment, make decisions, and take actions to achieve specific goals. Unlike traditional automation that follows fixed rules, agents can adapt their behavior based on context and learn from outcomes.
- **End-to-end process intelligence:** The ability for AI to understand, monitor, and optimize complete business processes that span multiple applications and departments—from initial customer inquiry through fulfillment, invoicing, and payment—with full awareness of business context, dependencies, and outcomes at each step.



100336enUS (26/02) © 2026 SAP SE or an SAP affiliate company. All rights reserved. No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies. Please see <https://www.sap.com/trademark> for additional trademark information and notices.

