


Beyond the AI Pilot

Six Questions for Turning Enterprise Data into Intelligent Action

A practical guide to building the data,
governance, and operating foundations
for production-ready AI.



A successful AI pilot can demonstrate technical potential, but production readiness begins with a different question: what measurable business outcome will it improve? To scale securely, reliably, and consistently, leaders must look beyond model performance alone and assess whether they have the right technical and operational foundation to support it. That requires establishing a trusted data foundation with rich business context, unified governance, and secure access across the environment.

Obstacles typically sit outside the model in the form of fragmented data, missing business context, uncertain accountability, inconsistent governance, or inability to quantify outcomes. As Rajeev Dasari, Global Director, SAP & Google Agentic Data Cloud Partnership & GTM at Google Cloud, puts it, “the model is just an enabler. A pilot proves the model can perform a task. Production proves that a business is ready to operationalize AI at scale.”


Katryn Cheng, VP of Product Marketing at SAP, describes the challenge as a gap between a staged pilot and the data of a real-world business environment. Pilot data is often prepared specifically for a tightly controlled sandbox, whereas a productive enterprise is hyperconnected across business functions, even when the underlying data layer is fragmented. As such, an AI system might perform adequately on staged data, but it may struggle when exposed to the disconnected data of day-to-day operations.

“Engineers try to model AI as closely as possible to the real world, but they’re thinking of it from their own perspectives, not from the actual people doing their jobs day-to-day,” she says.

Recent research reinforces the distinction between a successful demonstration and a real enterprise capability. In a March 2026 study of 51 successful enterprise AI deployments, Stanford Digital Economy Lab found that 77% of respondents said that the hardest implementation challenge was ‘invisible’ costs in areas like change management, data quality, and process redesign.¹

Before scaling AI, leaders must establish whether the use case aligns with a measurable outcome, whether it has access to the right information, and whether they have the governance needed to avoid unnecessary cost and risk. Taken together, the six questions we explore here test the six dimensions of production readiness:

1. **Business outcomes**
2. **Trusted data**
3. **Business context**
4. **Open and scalable architecture**
5. **Governance**
6. **Evidence and earned autonomy**



Question 1

What measurable business outcomes do we want to improve with AI?

In the rush to adopt AI, organizations can lose sight of the business reason for doing so. Leaders must therefore distinguish between an interesting technical demonstration, a use case capable of providing measurable value, and the operational requirements for scale.

Dasari recommends starting with several key questions: “Does the AI add value? Does that require an agent? What is the autonomy of that agent going to look like? All these things matter.” In other words, the model, architecture, and level of automation should all follow the business intent rather than define it.

Cheng similarly advises starting with an outcome the business already understands, such as reducing customer churn or



AI ambition still outpaces revenue impact


74%

Expect AI to grow revenue

20%

Report that it is already doing so

A January 2026 survey of **3,235 business leaders** found that 74% of organizations expect AI to grow revenue, but only 20% report that it is already doing so.²



Google Cloud



lowering inventory costs. That means defining the desired result, establishing a baseline, and deciding how the business will measure improvement before development begins.

When estimating ROI, leaders should also account for the costs of data engineering, governance, and ongoing maintenance to ensure that real-world expenses do not eclipse the business value. By setting a realistic target period for when the organization can expect to see value and what evidence will be required to justify further automation, leaders can also mitigate investment risk.

Dasari cites an invoice-processing initiative designed to prevent leakage caused by duplicate payments and mismatches between invoices and supplier contracts. The AI solution matches invoices received via multiple channels against contracts, checks purchase orders and payment status, and uses agents to route and validate the workflow.

Cheng points to an example from supply chain management, where the desired outcome is to reduce inventory costs and protect production when an unforeseen disruption impacts suppliers or routes. Such use cases are commercially meaningful when the business can state the intended result and track whether the intervention actually changed it.

Question 2

Does our AI have access to a current and complete view of the business?

AI can't support forward-looking decisions if it is grounded in stale or incomplete data, or data that's stripped of business meaning. Production AI must operate on information that reflects the business's current operational realities, especially when it needs to recommend time-sensitive actions.

For instance, use cases such as rerouting supply chains or sourcing alternative suppliers during a disruption require access to up-to-date data. By contrast, batch data is usually fine for running historical reports. As Dasari says, "Both are completely different use cases, but in supply chain, for example, we can't afford to have stale data."

However, for autonomous agents expected to judge, reason, and act within a business process, Cheng generally recommends that organizations rule out batch data in certain use cases. When the

Real-time supply chains deliver meaningful gains

↓ **20%**
Cost reduction

A May 2026 McKinsey study found that distributors who redesigned their supply chains with real-time data and AI achieved 20% cost reductions, along with meaningful gains in on-time delivery.³



system needs knowledge of promotions, inventory, budgets, and other current conditions that might change quickly and sometimes unexpectedly, the system needs real-time data. “If the agent makes a decision based on old data, then it could ruin the whole inventory for the next month or, in other agentic scenarios, do much worse,” she warns.

A complete business view rarely comes from a single system. For instance, autonomous agents may need live SAP, non-SAP and external data, with semantics and business context preserved when that information is brought together.

Dasari goes back to the example of supply chain management. “When you take the supply chain data and combine it with weather data, you can clearly know which parts of the world have weather events and how those might impact your suppliers.”

Cheng gives an example from retail planning, where SAP business data around products, sales, finance, and procurement can be combined with non-SAP and external signals such as marketing, demand, and market data. Together, these sources provide a more complete and current view of the business environment.

Question 3

Will our AI understand the meaning and context behind the data?

Fresh data tells AI what is happening now, but it does not necessarily explain why it matters and how different records relate. Beyond technical metadata and access to a transaction or record, AI needs semantic and business-process context to provide reliable reasoning. That might include information about entity relationships, process history, policies, or organizational structures.

AI pilots typically work within a limited scope and relatively small dataset. However, in a production environment, systems need business context that organizations often lose when extracting or copying data for analytics and AI.

“Data explains what happened, while business context explains why it matters,” says Dasari. For example, a purchase order alone will probably not mean much to an AI agent. The system might also



“Data explains what happened, while business context explains why it matters.”

Rajeev Dasari, Google Cloud

The semantic readiness gap

22%

Organizations have the data readiness required to scale advanced AI use cases

A May 2026 Strategy& survey of **50 CIOs in Germany** found that only 22% of organizations have the data readiness required to scale advanced AI use cases, suggesting that shared semantic models are vital for success.⁴

need to know who approved it and which contracts and policies it is connected to.

The problem with simply extracting and copying data is that it can separate data from the application-level business context. As Cheng says, “business context is in the relationships, the dependencies, and understanding of the business through data.” Rebuilding that meaning is hard because technical teams may not share the finance, merchandising, marketing, or supply-chain knowledge held by business experts.

SAP Business Data Cloud preserves meaning and semantics across a portfolio rather than within one isolated domain. That way, finance, workforce, contractor, supplier-network, and other data products can be connected to provide an end-to-end view that allows AI to reason over the entire operating model. According to Dasari, that is what allows AI to “make informed decisions and make the right recommendations, because it now understands the business operating model, not just the data behind it.”

Cheng emphasizes the role of SAP’s application-layer knowledge, which includes its data products, knowledge graph, metadata and business context. Because SAP understands how its tables, applications, and business processes relate, it can enrich data with relationships, lineage, metadata, and semantics. “We don’t need to infer context from the data, because we own the application layer, and so we know what the relationships are,” she says.

Google Cloud



Question 4

Can our organization provide access without multiplying data risk and cost?

AI systems often need access to data distributed across applications, platforms, and clouds. Yet the true cost of providing that access extends far beyond storage and compute. Teams also need to account for pipelines, reconciliation, semantic reconstruction, security administration, and ongoing governance.


“For every additional copy of enterprise data created in another place, the trust in that data has to be rebuilt,” Dasari explains. “Then there are latency considerations, reconciliation challenges, and the operational costs associated with all of that. It creates a multiplier effect.”

Cheng warns that every AI use case may require technical and business teams to rebuild context lost during data extraction

AI cost management becomes a key priority

The 6th Annual State of FinOps survey found that 98% of organizations now manage AI spend, up from 31% two years earlier. AI cost management also ranked as the leading skill set teams must develop.⁵

Organizations that manage AI spend





and replication, which increases the operational burden and delays business value realization. When incomplete or inaccurate information passes through an automated workflow, those weaknesses can compound across subsequent recommendations and autonomous actions.

The consequence goes beyond additional engineering expense. Every duplicate dataset can also create another version to secure, reconcile, update, and trace, which increases the possibility that AI acts on stale information or controls that no longer match the source system.

To reduce both cost and risk, Dasari recommends an architecture that minimizes unnecessary data movement while preserving governance and business context. “Rather than requiring organizations to move all their data to AI, the opportunity lies increasingly in bringing AI to trusted data where it resides, while preserving governance and business context,” he says.

Dasari notes that this approach is especially valuable in multi-cloud environments, where enterprise data may be distributed across applications, platforms and clouds. “AI can seamlessly reason over all of this data, across all the different clouds, making it relevant for customers to take advantage of the investments they’ve already made.”

Question 5

What should our AI agents be permitted to know, decide, and do?

AI can be unpredictable due to its inherently nondeterministic nature. In other words, it can make mistakes, especially if it lacks access to the right information. Because of this, agents must be properly governed, meaning that they should only have access to the data, business context, and systems they legitimately need to perform their assigned tasks.

The same governance principles that organizations have been using for decades still apply. As Dasari says, “AI agents should be governed with the same degree of rigor we expect for employees, with clear identity, authorization, least-privilege access, auditability and accountability appropriate to what the agent is permitted to do.”

What has changed with the rise of agentic AI is the speed and increasing scale at which agents are proliferating throughout the




Agent deployment is outpacing governance

30%

Organizations had reached level three on the Responsible AI (RAI) maturity model

McKinsey's March 2026 AI Trust Maturity Survey of around 500 organizations found that only 30% had reached level three on the [Responsible AI \(RAI\) maturity model](#).⁶



enterprise. Those same controls that organizations once depended on to monitor and govern interactions between employees and business applications and data must now be used when response times need to be measured in seconds rather than hours or days.

Cheng argues that governance must be established before agent deployment, rather than after. “Don’t deploy agentic AI without governance in place, because autonomous agents need to be watched, especially when they’re in the early stages of deployment,” she warns.

Her definition of agentic AI governance extends beyond technical controls to full visibility and human oversight. These controls include deciding whether agents inherit a user’s rights, removing destructive permissions that allow agents to permanently delete or modify critical data, requiring human approval for consequential workflows, and continuously monitoring for security and compliance.

Dasari recommends applying existing governance principles through a graduated operating process, in much the same way as an organization might when onboarding a new employee or promoting access rights to an existing one. “When we build agents, we basically have a human in the loop, and we want to validate it against the data and the business context,” he says. “If the business context changes, or the data changes, then you have to reevaluate the agent.”

Question 6

What evidence will justify further autonomy and investment in AI?

Most leaders recognize that AI is important, but many struggle to explain precisely why. Others might feel pressured to choose immediately between manual oversight and full automation, especially if they are in a highly regulated industry like finance or healthcare. Trust and autonomy should be earned through evidence, rather than being assumed. That means leaders need proof that a use case is improving the defined outcome, before expanding investment or autonomy.

Dasari returns to his earlier example of evidence-based progression through invoice reconciliation. “You can see I have validation at each and every step of the process before that agent becomes fully autonomous,” he says. “Each stage should demonstrate measurable improvement before advancing. Autonomy should be earned through evidence.”



Agent scale exposes oversight gaps

50%

Experienced growing reliability or quality problems as usage scaled

44%

Reported agents taking incorrect actions

62%

Currently have human-in-the-loop processes for agents

Among organizations piloting or deploying AI agents, half experienced growing reliability or quality problems as usage scaled, while 44% reported agents taking incorrect actions. Yet only 62% currently have human-in-the-loop processes for agents.⁷

Cheng goes further by challenging the assumption that a purpose-built agent is ready to carry out a broader role after proving itself. In her view, further autonomy may also require redesigning or creating a different agent, so she frames the question around whether a business has the right prerequisites. “If you want to judge whether your company is ready for autonomous agents, you need to judge your data foundation first to ensure your data has the context that agents need to make correct decisions,” she says. That rule applies to the deployment of any new agent.

Human-in-the-loop feedback is vital at every stage of this journey because, after all, you cannot hold a machine accountable. Dasari recommends treating traceability and auditability as requirements for further automation, with people validating agents in early production and remaining available to step in should conditions change. “If the context or the data changes, human-in-the-loop kicks in again to ensure you validate that the agent is still doing what it’s supposed to do,” he says.

Cheng describes human feedback as the most direct test of whether an agent is exercising sound judgment. For instance, in complex or consequential scenarios, agents can gather and synthesize data rapidly, but an experienced employee should still review the proposed plan before execution. This approach protects the business and helps seasoned users build trust in a system that impacts their responsibilities. “Autonomous doesn’t mean automatic. It’s about more autonomous agents helping humans,” she reminds us.



Conclusion

Enterprise readiness is a business discipline

Accurately and comprehensively answering the six questions we have presented here should be part of an integrated discipline, starting with a clear definition of the desired business outcomes. Production readiness therefore rests on the six connected dimensions we have explored above.

“Ask yourself, what are the business outcomes? What is the trusted data? What is the business context? Can we scale securely? Can AI act responsibly and transparently? And finally, has AI demonstrated measurable value to justify greater autonomy?” asks Dasari.

SAP and Google Cloud can support this progression by combining semantically rich SAP business data with Google Cloud data and its AI capabilities. That helps organizations preserve context and governance while moving priority use cases into production in less time.

The next step is simple: choose one priority AI use case, apply these six questions, identify the gaps, and let evidence—not ambition alone—determine when you’re ready to scale. SAP, in partnership with Google Cloud, can then help you build the data foundation needed to take it into production.


References

1. [The Enterprise AI Playbook](#), Stanford Digital Economy Lab, Stanford University, April 2026
2. [State of AI in the Enterprise: The Untapped Edge](#), Deloitte, January 2026
3. [AI is transforming distribution supply chains, turning margin pressure into a competitive advantage](#), McKinsey & Company, May 2026
4. [AI-driven change and the next big thing](#), Strategy&, May 2026
5. [State of FinOps 2026 Report](#), FinOps Foundation, February 2026
6. [State of AI Trust in 2026: Shifting to the Agentic Era](#), McKinsey & Company, March 2026
7. [The Value of AI](#), SAP & Oxford Economics, May 2026



About SAP

As a global leader in enterprise applications and business AI, SAP (NYSE:SAP) stands at the nexus of business and technology. For over 50 years, organizations have trusted SAP to bring out their best by uniting business-critical operations spanning finance, procurement, HR, supply chain, and customer experience. For more information, visit www.sap.com.



About Google Cloud

Google Cloud offers a powerful, optimized AI stack—including AI infrastructure, leading models like Gemini, data management capabilities, multicloud security solutions, developer tools and platform, as well as agents and applications — that enables organizations to transform their business for the Agentic Era. Customers in more than 200 countries and territories turn to Google Cloud as their trusted technology partner.

[Learn More](#)

Google Cloud





Expert led. Impact driven.

Studio is Informa TechTarget's global content studio offering brands an ROI rich tool kit: Deep industry expertise, first-party audience insights, an editorial approach to brand storytelling, and targeted distribution capabilities. Our trusted in-house content marketers help brands power insights-fueled content programs that nurture prospects and customers from discovery through to purchase, connecting brand to demand.

[Learn more](#)