

NEWS RELEASE

The next era of enterprise AI will be defined by ROI & trust

2025-12-10

The next era of enterprise AI will be defined by ROI & trust



0:00 / 1:51



By Heath Ramsey

The era of AI experimentation is over. Leaders now face a seemingly simple directive: turn AI investments into measurable value — fast.

Today, AI is everywhere, from service requests to IT incidents to onboarding and more. When technology becomes deeply embedded in business processes, expectations spike. In 2026, the winners will be the companies that make AI work at scale — not just in pilots, but across the entire enterprise.

Winners won't start with AI models; they'll start with processes

In 2026, leading companies won't begin with models or even with an AI tool — they'll begin with the work itself. Real impact will come from identifying a high-friction process, fixing it, validating the improvement under real conditions, and scaling the approach across the entire business. The shift isn't about isolated projects but solving a single meaningful problem so well that it becomes the template for enterprise transformation.

The method is straightforward: Choose a process that slows teams down, test an AI-driven approach under real-world pressure, run it through a single secure gateway, and use that success as a template for the next business objective. Each win strengthens a repeatable blueprint to reduce friction and scale across the business.

This is where ServiceNow, the AI Control Tower for business reinvention, comes into play. Instead of stitching models, systems, and policies together manually, ServiceNow unifies the entire AI value chain into one governed environment. Developers can define which models are approved for which tasks, what data and policies apply, and how every output is logged, traced, and audited.

When everything is visible and measurable in one place, trust becomes something leaders can operationalize.

ServiceNow's AI Control Tower keeps AI efforts governed and resilient as the landscape evolves. New AI services can be added, old ones retired, and policies can be updated without restarting governance from scratch. Centralized rules, automated routing, and consistent oversight eliminate one-off reviews and ad-hoc decisions that slow teams down.

In 2026, AI's impact will become especially clear in day-to-day operations. Take a benefits exception request, which today might bounce across three teams and linger for two days. With an AI agent operating through the AI Control Tower, the system gathers the necessary details, checks the request against policy, drafts a decision, and routes it for a quick human approval, often within the same day. And because that pattern is governed and reusable, the same approach quickly extends to payroll issues, leave requests, vendor changes, and more without rebuilding controls each time.

This is how AI stops being a set of disconnected pilots and becomes a reliable engine for business reinvention.

Trust at scale starts with action

Guardrails will make or break enterprise AI. When AI's actions are consistent, explainable, and governed,

enterprises can expand automation with confidence. Highly regulated industries, such as the public sector, healthcare, and financial institutions, must be able to trust that AI will behave predictably. As organizations increase agent autonomy, deterministic workflows – rules-based, predictable processes that ensure consistent outcomes – will become the bridge between experimentation and enterprise-wide adoption. Think of them as the “policy playbooks” AI must follow. Each step is explicit and auditable, ensuring decisions are driven by rules, not just probabilities. This is the foundation that allows automation to scale safely.

In 2026, organizations will treat deterministic workflows not as technical tooling, but as a core layer of digital governance and security. They form a rule-driven execution layer that makes automation faster, safer, and more efficient. Next-generation enterprise automation will blend two strengths: Deterministic steps for control and governance and probabilistic intelligence for interpretation and generation. Because of this, agentic playbooks will emerge as the new operating model for how work gets done — a vision already coming to life with the recent **ServiceNow Zurich platform release**.

Plus, with **ServiceNow’s recent announcement about its intent to acquire Veza**, a leader in identity security and governance, companies gain expanded visibility and control over identity governance. As organizations adopt more autonomous, agentic AI, this integration strengthens least-privilege access management and helps reduce cybersecurity risk.

Rollouts will decide whether AI sticks or stalls

A successful AI rollout isn’t just a single moment in time; it’s about an ongoing way of working. Enterprises must shift from only “projects” to full-scale programs, turning agent pilots into reliable, repeatable success stories. AI demands continuous testing, always-on monitoring, and full visibility across systems and models. When teams have a single view of all AI activity, they can spot behavior patterns, identify hotspots, correct failures, and surface emerging risks. That visibility is what keeps AI reliable — not just on day one, but every day after.

A great example of how to do this right is **IBM**. IBM’s CIO organization operationalized AI at scale by instrumenting workflows end-to-end, pressure-testing before launching, and running continuous monitoring across thousands of employee support requests. The results: Faster resolutions, fewer escalations, and a repeatable pattern they can efficiently apply to new processes.

The path forward is simple, but not easy: Fix real work. Build real trust. Scale real wins.

In 2026, AI isn’t a bet. It’s the way business runs.