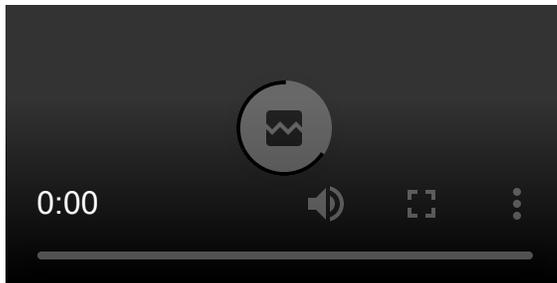


NVIDIA GTC 2026: Governing the Autonomous Workforce

2026-03-16



By: Joe Davis, Executive Vice President of AI Platform and Product Engineering

There's no shortage of bold claims about AI reshaping the software landscape. As models become more capable and agents more autonomous, AI is moving from experimentation to production at unprecedented speed.

What's becoming clear is that generating answers is only the starting point, and success at enterprise scale requires platforms that can operationalize, govern, and connect AI to real workflows. At **NVIDIA GTC 2026**, that reality came into focus.

ServiceNow and NVIDIA have been building toward this reality with a deep existing collaboration to unite intelligent workflows and models so that trusted AI can be deployed across industries.

Kari Briski, vice president, Generative AI Software at NVIDIA, put it well: "Our collaboration with ServiceNow spans the full AI lifecycle — from training and deployment on NVIDIA accelerated computing to governance and operational efficiency. With AI agents built on open models for the ServiceNow AI Platform, we're giving enterprises

everything they need to move AI 'thinking' to AI 'doing' at scale."

Building the foundation for a digital workforce

As AI agents move from isolated assistants to teams of autonomous digital workers driving critical enterprise workflows, new challenges emerge around coordination and governance.

Today at GTC, we're demonstrating how the recently launched **ServiceNow Autonomous Workforce of AI Specialists**, built on the ServiceNow AI Platform, can use the **NVIDIA Agent Toolkit**, including the **NVIDIA AI-Q blueprint**. These long-running agents augment human teams and leverage NVIDIA's accelerated infrastructure and AI frameworks with closed and open models, including April models and **NVIDIA Nemotron** running on **NVIDIA Blackwell** AI infrastructure.

NVIDIA powers modern AI systems, while ServiceNow provides the enterprise platform where those capabilities are combined with enterprise context and intelligent workflows. For example, Level 1 Service Desk AI Specialists, operating like teammates, can analyze incoming support tickets, investigate root causes using knowledge and historical data, and execute remediation — completing not just one task, like most current agents and bots, but coordinating entire workflows while enforcing governance and operational controls.

Together, ServiceNow and NVIDIA enable organizations to move beyond AI experiments to real business impact with governed, enterprise-ready AI execution.

From infrastructure to enterprise control

At NVIDIA GTC DC last October, ServiceNow and NVIDIA **announced** plans to integrate intelligent ServiceNow workflows with the NVIDIA Enterprise AI Factory validated design, reimagining how AI infrastructure connects to enterprise operations. At GTC 2026, that vision takes a meaningful step forward, representing a shift from isolated AI infrastructure to integrated enterprise AI architecture.

We're previewing a new integration between the **NVIDIA Enterprise AI Factory**— a full-stack validated design for enterprises to build and deploy their on-premise AI factories — and **ServiceNow AI Control Tower** — where models, agents, and prompts from any system are governed, monitored, and aligned to enterprise policy.

NVIDIA's Enterprise AI Factory provides a validated design for purpose-built environments to train and deploy AI workloads at scale. But as these environments expand across on-premises, sovereign, and hybrid deployments, enterprises face a new challenge: governing AI within operational and regulatory boundaries. ServiceNow AI Control Tower helps enable that oversight, extending observability and governance to any enterprise AI environment.

This early look signals the direction forward: AI infrastructure powered by NVIDIA, orchestrated, governed, and actioned through ServiceNow, scaling innovation without compromising trust.

Establishing enterprise-ready benchmarks for voice and multimodal AI

Our collaboration with NVIDIA also continues to advance how AI is evaluated and deployed as human interactions with the technology expand into text, voice, and multimodal workflows.

Together with NVIDIA, we're initiating an enterprise-focused benchmarking framework for voice and multimodal AI, leveraging NVIDIA Nemotron speech open-source models natively embedded within the ServiceNow AutoEval suite.

Unlike traditional academic benchmarks or vendor-specific scorecards, this is designed to simulate enterprise conditions, including conversational reasoning with voice agents — such as multi-turn employee requests — or areas like document intelligence and workflow execution under real-world policy constraints.

This will ultimately establish transparent, reproducible, enterprise-relevant evaluation standards that organizations can use for model selection, pre-deployment validation, and ongoing governance.

An architectural inflection point

GTC marks the next chapter in the ongoing collaboration between ServiceNow and NVIDIA, advancing a governed agentic workforce built on NVIDIA infrastructure, expanding AI Factory and AI Control Tower integration, and establishing enterprise-ready benchmarks for voice and multimodal AI.

At ServiceNow's **Knowledge 2026** conference in May, we'll share a deeper view with NVIDIA on the AI Control Tower and AI Factory integration, including expanded use cases and go-to-market execution designed to help organizations build, govern, and scale AI with confidence.

AI may reshape software, but at enterprise scale, it requires cross-system orchestration, operational controls, and accountability. The next phase of innovation isn't just about deploying AI; it's about breaking down silos and governing AI together.