



NVIDIA Ignites the Next Industrial Revolution in Knowledge Work With Open Agent Development Platform

NVIDIA Agent Toolkit Equips Enterprises to Build and Run AI Agents

News Summary:

- NVIDIA Agent Toolkit includes NVIDIA OpenShell open source runtime for building self-evolving agents and claws with more safety and security.
- Built with LangChain, the open source NVIDIA AI-Q Blueprint for agentic search tops DeepResearch Bench accuracy leaderboards — using a hybrid approach with both frontier and open models that can cut query costs in half.
- Leading software platforms including Adobe, Atlassian, Amdocs, Box, Cadence, Cisco, Cohesity, CrowdStrike, Dassault Systèmes, IQVIA, Red Hat, SAP, Salesforce, Siemens, ServiceNow and Synopsys are advancing enterprise and physical AI agents with NVIDIA Agent Toolkit software.

GTC—NVIDIA today announced that it is teaming with partners to ignite the next era of AI with open source software for autonomous, self-evolving enterprise AI agents — increasing agent safety, security and efficiency — to speed a generational shift in software and knowledge work.

NVIDIA Agent Toolkit provides open source models and software for enterprises and developers building tools that scale productivity by autonomously determining how to complete assigned tasks. It now includes [NVIDIA OpenShell](#)™ — an open source runtime that enforces policy-based security, network and privacy guardrails that make autonomous agents, or claws, safer to deploy.

“Claude Code and OpenClaw have sparked the agent inflection point — extending AI beyond generation and reasoning into action,” said Jensen Huang, founder and CEO of NVIDIA. “Employees will be supercharged by teams of frontier, specialized and custom-built agents they deploy and manage. The enterprise software industry will evolve into specialized agentic platforms, and the IT industry is on the brink of its next great expansion.”

NVIDIA Agent Toolkit Speeds Community and Enterprise Software Evolution

NVIDIA Agent Toolkit includes open models like [NVIDIA Nemotron](#)™, open agents like NVIDIA AI-Q, open skills like [NVIDIA cuOpt](#)™ and open runtimes like [OpenShell](#). Developers can use Agent Toolkit software for creating specialized AI agents that can act autonomously, while using and building other software to complete tasks.

Using the [NVIDIA AI-Q](#) open agent blueprint, developers can create custom AI agents that can perceive, reason and act on enterprise knowledge, automatically choosing the right data sources and depth of analysis to deliver precise, context-aware answers. A built-in evaluation system explains how each AI answer is produced.

The AI-Q hybrid architecture uses frontier models for orchestration and NVIDIA Nemotron open models for research, which can cut query costs by more than 50% while providing world-class accuracy. NVIDIA used the AI-Q Blueprint to develop the top-ranking AI agent on the [DeepResearch Bench](#) and [DeepResearch Bench II](#) leaderboards.

NVIDIA OpenShell gives [autonomous agents](#) the access they need to be productive while enforcing policy-based security, network and privacy guardrails. NVIDIA is collaborating with security providers including Cisco, CrowdStrike, Google, Microsoft Security and TrendAI to build OpenShell compatibility with their cyber- and AI-security tools.

Agent engineering company LangChain — which provides open source frameworks that have been downloaded over 1 billion times — is working with NVIDIA to integrate Agent Toolkit, including AI-Q, OpenShell and Nemotron open models, datasets and tools, into the LangChain deep agent library for developing advanced, accurate enterprise AI agents and running them at scale.

Enterprise Software Platforms Scale AI Agent Skills With NVIDIA

Software leaders across industries are working with NVIDIA Agent Toolkit software including OpenShell and NVIDIA Nemotron open models to expand the capabilities of AI agents in their applications and platforms.

- [Adobe](#) will adopt Agent Toolkit software as foundations for running hybrid, long-running Adobe creativity, productivity and marketing agents in a personalized, more secure and cost-efficient environment.
- Amdocs is using NVIDIA AI-Q and Nemotron to power its Cognitive Core agent platform that continuously monitors customer interactions and billing data to proactively identify and resolve issues before customers are impacted.

- [Atlassian](#) is working with NVIDIA Agent Toolkit software and NVIDIA OpenShell as it continues to evolve its Rovo AI agentic strategy and AI-powered system of work for tools like Jira and Confluence.
- [Box](#) is using NVIDIA Agent Toolkit to enable enterprise agents using the Box file system to securely and reliably execute long-running business processes.
- [Cadence](#) will leverage the Agent Toolkit and Nemotron with the Cadence ChipStack AI SuperAgent to help engineers design and verify higher-quality, more complex semiconductor designs.
- [Cisco](#) AI Defense will provide AI security protection for OpenShell, adding controls and guardrails to govern agent and claw actions.
- [Cohesity](#) is adding support for NVIDIA OpenShell and expanding its Gaia AI platform to support more advanced agentic workflows with NVIDIA AI-Q.
- [CrowdStrike](#) unveiled a Secure-by-Design AI Blueprint that embeds Falcon platform protection directly into NVIDIA AI agent architectures — including agents built on NVIDIA AI-Q and OpenShell — and is advancing agentic managed detection and response by using NVIDIA Nemotron reasoning models and [NeMo Data Designer](#) to power AI agents for investigative workflows.
- [Dassault Systèmes](#) is exploring NVIDIA Agent Toolkit software and NVIDIA Nemotron open models for its role-based AI agents, called Virtual Companions, on the 3DEXPERIENCE agentic platform.
- [IQVIA](#) is integrating NVIDIA Nemotron and other Agent Toolkit software with IQVIA.ai, a unified agentic AI platform, designed to help life sciences organizations work more efficiently, make better decisions and scale AI responsibly across clinical, commercial and real-world operations. IQVIA has deployed more than 150 agents across internal teams and client environments, including 19 of the top 20 pharma companies.
- [Palantir](#) is working with NVIDIA, using Nemotron for developing AI agents that run on the Palantir sovereign AI Operating System Reference Architecture.
- [Red Hat](#) is integrating Agent Toolkit software into Red Hat AI Factory with NVIDIA, providing an enterprise-ready platform for building more secure, autonomous agents.
- [Salesforce](#) is working with NVIDIA Agent Toolkit software including NVIDIA Nemotron models, enabling customers to build, customize and deploy AI agents using Agentforce for service, sales and marketing tasks. This collaboration introduces a reference architecture where employees can use Slack as the primary conversational interface and orchestration layer for Agentforce agents — powered by NVIDIA AI infrastructure — that participate directly in business workflows and can pull from data stores in both on-premises and cloud environments.
- [SAP](#) is using open NVIDIA Agent Toolkit software including NVIDIA NeMo™ for enabling AI agents through Joule Studio on SAP Business Technology Platform, enabling customers and partners to design agents tailored to their own business needs.
- [Siemens](#) is launching the [Fuse EDA AI Agent](#) which uses NVIDIA Nemotron to autonomously orchestrate domain-scoped workflows across Siemens' electronic design automation portfolio for the semiconductor and printed circuit board industries, from design conception through manufacturing sign-off, increasing engineering efficiency and design quality.
- [ServiceNow](#)'s Autonomous Workforce of AI Specialists are built on the ServiceNow AI Platform and leverage NVIDIA Agent Toolkit software, the NVIDIA AI-Q Blueprint, and a combination of closed and open models, including NVIDIA Nemotron and ServiceNow Apriel models.
- [Synopsys](#) is building a multi-agent framework, powered by its AgentEngineer technology, for semiconductor and systems design using NVIDIA Nemotron and Nemo Agent Toolkit.

Developers can explore NVIDIA Agent Toolkit and OpenShell on build.nvidia.com today, running on inference providers and NVIDIA Cloud Partners including Baseten, Bitdeer AI, CoreWeave, DeepInfra, DigitalOcean, GMI Cloud, Fireworks, Lightning, Together AI and Vultr.

Developers can also use OpenShell with LangChain, and download it from [GitHub](#) to run locally on dedicated [NVIDIA GeForce RTX™ PCs and laptops](#) or [NVIDIA RTX™-powered workstations](#), as well as [NVIDIA DGX Station™](#) and [NVIDIA DGX Spark™](#) supercomputers from [Altos Computing](#), [ASUS](#), [Dell Technologies](#), [GIGABYTE](#), [HP](#), [Lenovo](#), [MSI](#) and [Supermicro](#).

Enterprises can also build and run AI agents on AI factory infrastructure from cloud service providers including Amazon Web Services (AWS), Google Cloud, Microsoft Azure, Microsoft Security and Oracle Cloud Infrastructure, and on servers from Cisco, Dell Technologies, [HPE](#), Lenovo and Supermicro.

Watch the [GTC keynote](#) from Huang and explore [sessions](#).

About NVIDIA

[NVIDIA](#) (NASDAQ: NVDA) is the world leader in AI and accelerated computing.

Certain statements in this press release including, but not limited to, statements as to: Claude Code and OpenClaw extending AI beyond generation and reasoning into action; employees being supercharged by teams of frontier, specialized and custom-built agents they deploy and manage; the enterprise software industry evolving into specialized agentic platforms, and the IT industry is on the brink of its next great expansion; the benefits, impact, performance, and availability of NVIDIA's products, services, and technologies; expectations with respect to NVIDIA's third party arrangements, including

with its collaborators and partners; expectations with respect to technology developments; expectations with respect to AI and related industries; and other statements that are not historical facts are forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, which are subject to the “safe harbor” created by those sections based on management’s beliefs and assumptions and on information currently available to management and are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic and political conditions; NVIDIA’s reliance on third parties to manufacture, assemble, package and test NVIDIA’s products; the impact of technological development and competition; development of new products and technologies or enhancements to NVIDIA’s existing product and technologies; market acceptance of NVIDIA’s products or NVIDIA’s partners’ products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of NVIDIA’s products or technologies when integrated into systems; NVIDIA’s ability to realize the potential benefits of business investments or acquisitions; and changes in applicable laws and regulations, as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company’s website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

Many of the products and features described herein remain in various stages and will be offered on a when-and-if-available basis. The statements above are not intended to be, and should not be interpreted as a commitment, promise, or legal obligation, and the development, release, and timing of any features or functionalities described for our products is subject to change and remains at the sole discretion of NVIDIA. NVIDIA will have no liability for failure to deliver or delay in the delivery of any of the products, features or functions set forth herein.

© 2026 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, cuOpt, DGX Spark, DGX Station, GeForce RTX, NVIDIA NeMo, NVIDIA Nemotron, NVIDIA OpenShell and NVIDIA RTX are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

Corporate Communications
NVIDIA Corporation
press@nvidia.com