



Enterprise Software Leaders Build AI Agents With NVIDIA

News Summary:

- New NVIDIA Agent Toolkit software -- including latest NVIDIA NemoClaw blueprints, Nemotron models, OpenShell secure runtime and CUDA-X libraries with agent skills -- delivers open source foundations for enterprise development.
- Cadence, Dassault Systèmes, Siemens and Synopsys are among the first to use NVIDIA NemoClaw to build autonomous AI engineers working as digital coworkers to execute simulation and verification workflows -- compressing weeks of engineering work into hours.
- New NVIDIA Nemotron 3 Ultra is a smaller, faster open model built for long-running agents, delivering 5x faster inference and up to 30% lower cost for complex agentic tasks.
- CrowdStrike and Palantir are transforming cybersecurity and operational decision-making with long-running AI agents powered by Nemotron open models, enabling teams to analyze data faster and streamline operations across complex environments.
- NVIDIA and Microsoft are collaborating to deliver a native Windows experience for personal agents with new security primitives and NVIDIA OpenShell, while Canonical and Red Hat are integrating NVIDIA OpenShell as a secure, open source runtime with policy and privacy controls for agents, extending support across PCs, data centers and clouds.
- NVIDIA CUDA-X libraries -- including cuDF, cuOpt, AI-Q, NeMo, PhysicsNeMo and CUDA-Q -- are now accessible to AI agents as domain-specific skills to expand agent capabilities.

NVIDIA GTC Taipei--NVIDIA today announced new software, open source models and partnerships with the world's leading software platform providers to build autonomous AI agents for industries and enterprises across engineering, healthcare, software development and business operations.

Leading software companies are using NVIDIA Agent Toolkit software to build secure, long-running AI agents that act as digital coworkers. These autonomous agents start with a model. Then, they require a software layer called a harness to turn the model into an agent with functions like orchestration, context, memory, tool use and security.

NVIDIA Agent Toolkit software equips enterprises to build agents that can work alongside employees at scale. [NVIDIA Nemotron](#)™ open models and [NVIDIA NemoClaw](#)™ blueprints connect popular harnesses; the [NVIDIA OpenShell](#)™ secure runtime sets policy and privacy controls; and agents can now tap into [NVIDIA CUDA-X](#)™ libraries as skills.

"The world's software leaders are bringing AI agents into the systems where work gets done -- showing how AI coworkers help employees think faster and execute complex tasks to solve bigger problems," said Jensen Huang, founder and CEO of NVIDIA. "NVIDIA NemoClaw provides enterprise software developers with the open building blocks to create more secure, long-running AI coworkers that amplify human expertise as they reshape how work gets done."

Design and Simulation Leaders Build Autonomous AI Engineers With NVIDIA NemoClaw

In semiconductor and industrial engineering, simulation and verification are among the most time-intensive workloads, requiring teams to execute complex, repetitive workflows across days or weeks before a design can move forward.

Cadence, Dassault Systèmes, Siemens, Synopsys, [Flexcompute](#), Luminary, Neural Concept, [nTop](#), [P-1 AI](#), PhysicsX, SimScale and [Synera](#) are among the first to build autonomous AI engineers to work alongside employees, using NVIDIA NemoClaw. By delegating these tasks to always-on autonomous AI engineers, organizations can compress those weeks of engineering cycles into hours and redirect human expertise toward the work that demands it most.

[Cadence](#) is using NVIDIA OpenShell to secure its [ChipStack AI Super Agent](#), a fully autonomous AI engineer that executes chip design and verification. NVIDIA is the first customer using ChipStack to autonomously verify its chip designs.

[Dassault Systèmes](#) is using NVIDIA NemoClaw and OpenShell to productize the 3DEXPERIENCE agentic platform for long-running, autonomous agents across design, simulation and manufacturing operations.

[Siemens](#) is integrating NVIDIA NemoClaw and OpenShell into Fuse EDA AI Agent, a purpose-built autonomous agent that plans and orchestrates multi-tool workflows across semiconductor, 3D integrated circuit and printed circuit board system design.

[Synopsys](#) is working with NVIDIA to build always-on autonomous AI engineers for chip design, with a focus on achieving full workflow autonomy.

Also announced today, [Foxconn](#) is piloting NVIDIA NemoClaw to power its Nurabot and CoDoctor platforms, using teams of specialized AI agents to support clinical reasoning, documentation and care coordination. [Foxconn](#) is also using the [NVIDIA](#)

[FOX](#) and NemoClaw blueprints to build MoMClaw, a factory operations agent that connects sensor and machine data with AI agents to deliver real-time insights and action plans with NVIDIA OpenShell privacy controls and safety guardrails.

NVIDIA Unveils New Nemotron Open Models to Power Long-Running Agents

Fueling the intelligence of autonomous agents, NVIDIA today unveiled new open models and datasets for always-on agents, developed with contributions from the [NVIDIA Nemotron Coalition](#).

NVIDIA Nemotron 3 Ultra is a 550-billion-parameter mixture-of-experts model that delivers frontier-level intelligence to long-running agents across coding, research and enterprise workflows. With up to 5x faster inference and up to 30% lower cost compared with open frontier models in its class, Ultra enables agents to complete tasks faster and at lower cost.

Nemotron 3 Ultra is post-trained for leading agent platforms and harnesses -- the orchestration frameworks enterprises use to deploy and coordinate agents -- including Hermes Agent, LangChain Deep Agents, OpenClaw, OpenHands and OpenCode.

In addition, new Nemotron models for safety and speech recognition further expand the model family's capabilities for building efficient, specialized enterprise agents.

NVIDIA Nemotron models are enabling a new class of long-running AI agents across enterprise platforms from companies including CrowdStrike and Palantir. These agents help teams analyze complex data, coordinate tasks and streamline operations across cybersecurity and enterprise environments.

CrowdStrike is using NVIDIA Nemotron models for its specialized agents that continuously identify, prioritize and remediate vulnerabilities and policy misconfigurations, helping stop adversaries faster while reducing the operational burden on security teams.

Palantir is integrating NVIDIA Nemotron models into its AI FDE (Forward Deployed Engineer) platform to autonomously execute complex tasks, enabling continuous learning from agent interactions to build domain-specific, air-gapped enterprise systems.

Major Software Platforms Integrate Secure Agent Runtime With NVIDIA OpenShell

Autonomous agents that write code, generate sub-agents and remember context across sessions can access local files, learn new tools and execute advanced workflows with increasing independence. The more capable agents become, the more important it is to have necessary guardrails for the agents to operate within. The critical layer is a runtime with adjustable privacy and security controls that make autonomous agents safer to deploy at scale.

[NVIDIA and Microsoft are partnering](#) across new Windows security primitives and the NVIDIA OpenShell runtime to ensure agents run safely and under full user control.

The new Windows primitives deliver identity, containment, policy and end-to-end security capabilities to build and run agents natively. NVIDIA OpenShell builds on these primitives to provide additional policy capabilities and will intelligently route queries to local models based on the user's privacy policies, as well as disguise personal information in queries sent to cloud models.

[Canonical](#) will integrate OpenShell with Ubuntu through supported snaps and rocks, aka OCI-compliant containers, to run autonomous agents on enterprise servers worldwide.

[Red Hat](#) is integrating OpenShell into its full-stack Red Hat AI platform to maintain oversight and policy at the infrastructure level. The company is also making key contributions to the OpenShell upstream open source project to help standardize how agents are managed on enterprise platforms.

Today's announcements build on recent integrations by [SAP](#), which is embedding OpenShell into Joule Studio runtime -- part of SAP Business AI Platform for enterprise AI agents -- and [ServiceNow](#), which secured Project Arc, ServiceNow's enterprise autonomous desktop agent, with OpenShell to add policy-based management for enterprise safety.

OpenShell runs in on-premises, hybrid and enterprise cloud environments, local devices such as [NVIDIA RTX Spark™](#), NVIDIA DGX Spark™ and GB10 systems from system providers, as well as [NVIDIA DGX Station™ for Windows](#) and NVIDIA DGX Station GB300 systems from NVIDIA partners.

NVIDIA CUDA-X Libraries Available as Skills for Autonomous Agents

NVIDIA CUDA-X libraries are now accessible to AI agents with [domain-specific skills](#), giving AI agents the ability to more easily use specialized capabilities to tackle the biggest challenges in science, industry and enterprise. Examples include:

- **NVIDIA cuDF** accelerates data processing and analytics over massive structured datasets, enabling agents to quickly reason over enterprise data.
- **NVIDIA cuOpt™** helps agents solve complex routing, scheduling, resource allocation, supply chain and decision optimization problems in real time.
- **NVIDIA AI-Q** gives agents intelligent routing, persistent context and built-in evaluation for enterprise research and knowledge workflows.

- **NVIDIA NeMo™** accelerates agent optimization, evaluation and governance through prompts, skills, model routing and model customization for specialized domains and regions.
- **NVIDIA PhysicsNeMo™** empowers agents to build and benchmark high-fidelity AI physics models for complex scientific and engineering simulations.
- **NVIDIA CUDA-Q™** enables AI agents to streamline installation, generate and test quantum programs, simulate quantum systems and orchestrate quantum applications.

NVIDIA [also released](#) a major collection of open source physical AI libraries, skills, models and frameworks, enabling [AI agents](#) and developers to stand up workflows that accelerate development of robotics, autonomous vehicles and industrial systems.

Availability

NVIDIA NemoClaw is [available now](#). NVIDIA OpenShell is available in [early preview](#).

NVIDIA Nemotron 3 Ultra is expected to be available on June 4 via Hugging Face, ModelScope, OpenRouter and [build.nvidia.com](#) as NVIDIA NIM™ microservices, as well as through a broad ecosystem of [NVIDIA Cloud Partners](#), inference platforms and cloud service providers.

The verified NVIDIA agent skills are available in the Claude Code plug-in marketplace, as well as Hermes Skills Hub.

Watch Huang's [keynote](#) and learn more at [NVIDIA GTC Taipei](#).

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: the world's software leaders bringing AI agents into the systems where work gets done — showing how AI coworkers help employees think faster and execute complex tasks to solve bigger problems; expectations with respect to growth, performance, availability, and benefits of NVIDIA's products, services and technologies, and related trends and drivers; expectations with respect to NVIDIA's third party arrangements, including with its collaborators and partners; expectations with respect to technology developments, and related trends and drivers; projected market growth and trends; 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 products 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, CUDA-X, cuOpt, NeMo, NemoClaw, Nemotron, NVIDIA DGX Spark, NVIDIA DGX Station, NVIDIA OpenShell, NVIDIA PhysicsNeMo and NVIDIA RTX Spark 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.

Corporate Communications
NVIDIA Corporation
press@nvidia.com