

TANIUM

Tanium for NVIDIA AI Infrastructure

Bring DGX Spark, RTX Spark, and DGX Station into the same real-time view as the rest of your endpoints — so your AI endpoints don't become your next blind spot.

AI Is Moving Fast. Your Visibility Has to Keep Up.

The future of AI compute introduced new endpoints to your environment faster than inventories can keep up — and each one that goes unseen widens your attack surface. NVIDIA's DGX Spark, RTX Spark, and DGX Station put frontier-scale AI compute on the desktop and at the deskside — a new class of endpoint that IT needs to see, secure, and manage like everything else.

You can't protect what you can't see

New system will appear faster than inventories update.

Data is siloed

Separate tools give IT and Security different answers about the same system.

Drivers and firmware drift

Inconsistent versions across clusters cause failures and open gaps.

Patching risks uptime

Long-running AI jobs make maintenance windows hard to find.

Tanium and NVIDIA: Enterprise Manageability for the AI Compute Era

Support for NVIDIA powered AI Endpoints — Tanium now supports NVIDIA's newest AI hardware — DGX Spark, RTX Spark, RTX Spark Pro, and DGX Station — with the same agent, real-time query, and policy-layer control used across the rest of the enterprise fleet. No separate tools, no custom builds.

Tanium is the enterprise endpoint platform named in NVIDIA's official DGX Spark Enterprise Manageability Guide — published inside the DGX Spark User Guide alongside Ansible, Puppet, and Chef. Where those tools stop at configuration push, Tanium adds the real-time query, natural-language investigation, and policy-layer runtime control enterprises actually need to run AI infrastructure with the same confidence as the rest of their fleet.

See It. Fix It. Run It.

One agent, one console, and one real-time data model across every NVIDIA AI endpoint in the enterprise:

Know every AI endpoint

Real-time inventory of hardware, drivers, and software across your AI estate, including unmanaged assets.

Keep workloads running

Update drivers and firmware with control over timing, so AI jobs stay up while systems stay current.

Close exposure fast

Find vulnerable and misconfigured systems, then patch and remediate at scale without waiting on scans.

Act with one platform

IT and Security work from the same live data and the same control plane.

- **DGX Spark and DGX Station** — Real-time patching, hardware and software inventory, configuration compliance, vulnerability visibility, and process management for desktop and deskside AI supercomputers — no custom build required. IT gets the same live answers for its AI systems that it already gets for the rest of the fleet.
- **RTX Spark and RTX Spark Pro** — Full management for AI-ready OEM laptops and workstations: OS and third-party patching, software deployment, configuration and compliance policy, and real-time query — the full Tanium endpoint management toolset, not a reduced feature set.

Tanium Atlas: Agentic AI for the AI Fleet

Tanium Atlas, now generally available in Tanium Cloud, gives a single IT or security operator the full weight of the Tanium Autonomous IT Platform. For teams managing NVIDIA AI powered infrastructure, that means:

- **Ask in Plain Language** — “Can you correlate GPU activity with AI Agent activity in a time series graph in a 24hr look back? I want to see all available metrics.”

Five Top Use Cases for Tanium Managed AI Endpoints

Five live demonstrations, each answering a question enterprises are already asking about their AI fleets:

01

AI Model Visibility

Get real-time views of every endpoint and AI model, with health state and resource footprint for each.

TANIUM DIFFERENCE

AI models visible on the same real-time fleet platform

02

AI Agent Lifecycle Management

Know which agents run, where, and what was done. Pause them without taking the entire system offline.

TANIUM DIFFERENCE

Manage agents like any other enterprise asset

03

AI Agent Containment

Control which external services agents can reach, without restarting the agent (for example, allow or block access to Hugging Face).

TANIUM DIFFERENCE

Live activity evidence and control in one platform

04

AI Stack Health & Drift

Find unsupported builds and driver drift to prevent failures and troubleshoot root cause faster.

TANIUM DIFFERENCE

Proven endpoint management applied to the AI stack

05

Model Economics & GPU Demand

Connect workload activity with GPU demand to choose models on evidence.

TANIUM DIFFERENCE

Endpoint, agent, model and hardware in one view

Why Tanium: Supporting the Next Evolution of the AI Endpoint

- **Real-Time at Any Scale** — Tanium's patented communications architecture returns live answers from millions of endpoints in seconds and acts on them just as fast — trusted across more than 36 million endpoints worldwide.
- **Industry-Recognized Leadership** — Named a Leader in the 2026 Gartner® Magic Quadrant™ for Endpoint Management Tools.

SUMMARY

A Foundation for the Enterprise AI Endpoint

Tanium brings enterprise-grade visibility, control, and real-time action to NVIDIA's AI infrastructure — DGX Spark, RTX Spark, and DGX Station — with Tanium Atlas turning live endpoint data into governed, agentic action. As enterprises scale their AI factories, Tanium gives them the same confidence in the AI endpoint they already expect everywhere else.

TANIUM

Visit us at tanium.com and follow us on LinkedIn and X.

Autonomous IT. **Unstoppable Business.**

© Tanium 2026