

PaletteAI VerteX

The mission-ready platform to securely design, deploy and manage compliant AI environments at scale.

The situation

Enterprise infrastructure is undergoing its fastest transformation in decades.

AI has become the defining force reshaping infrastructure, software ecosystems, and applications — all at once.

As organizations rush to harness its potential, the stakes are even higher for government and highly regulated industries, where security, compliance, and data integrity are paramount. The rapid adoption of AI introduces new risks, from data exposure to model tampering, making trusted, standards-based protection like FIPS 140-3 encryption essential from the start.

To maintain competitive advantage, enterprises are racing to build AI-ready infrastructure, taking advantage of exponential advances in GPU performance and new technologies like DPUs, which have emerged as critical enablers of Zero-Trust computing, secure data, and workload isolation.

Riding this AI wave is easier said than done. ROI is hard to come by. Overprovisioning and underutilization mean scarce, valuable resources are sitting idle, resulting in millions in wasted enterprise spend. In regulated industries, unmanaged AI growth can introduce risks around operational security.

Friction between teams is also growing.

AI and ML teams of data scientists and developers push for innovation, speed, and access to the latest tools. The AI application ecosystem is already 2.5x larger than the cloud-native ecosystem and evolving rapidly.

At the same time, platform teams demand control, compliance, and cost efficiency, trying to bring order to chaos, but they struggle to manage the sprawl of disconnected environments and the rise of “shadow AI” which often operates outside approved governance frameworks.

What’s the answer?

Enterprises need a secure, unified platform that bridges these competing needs: a solution that provides freedom and flexibility without sacrificing governance, FIPS-140-3-grade security, or ROI.

The solution

With PaletteAI VerteX, enterprises can build repeatable, compliant, and trusted AI workloads that accelerate innovation while maintaining operational discipline and meeting stringent security standards such as FIPS 140-3.

PaletteAI VerteX gives platform teams and AI practitioners a shared environment to design, deploy, and manage AI workloads from the data center to the edge with speed, confidence, and verified protection.

It includes distinct interfaces for platform teams and practitioners. Platform teams make available secure, compliant cluster templates in PaletteAI VerteX for AI teams to use. These clusters are built on standardized reference designs from NVIDIA and include all the essential infrastructure layers: operating system, Kubernetes, networking, and storage, providing a consistent and reliable foundation for AI workloads.

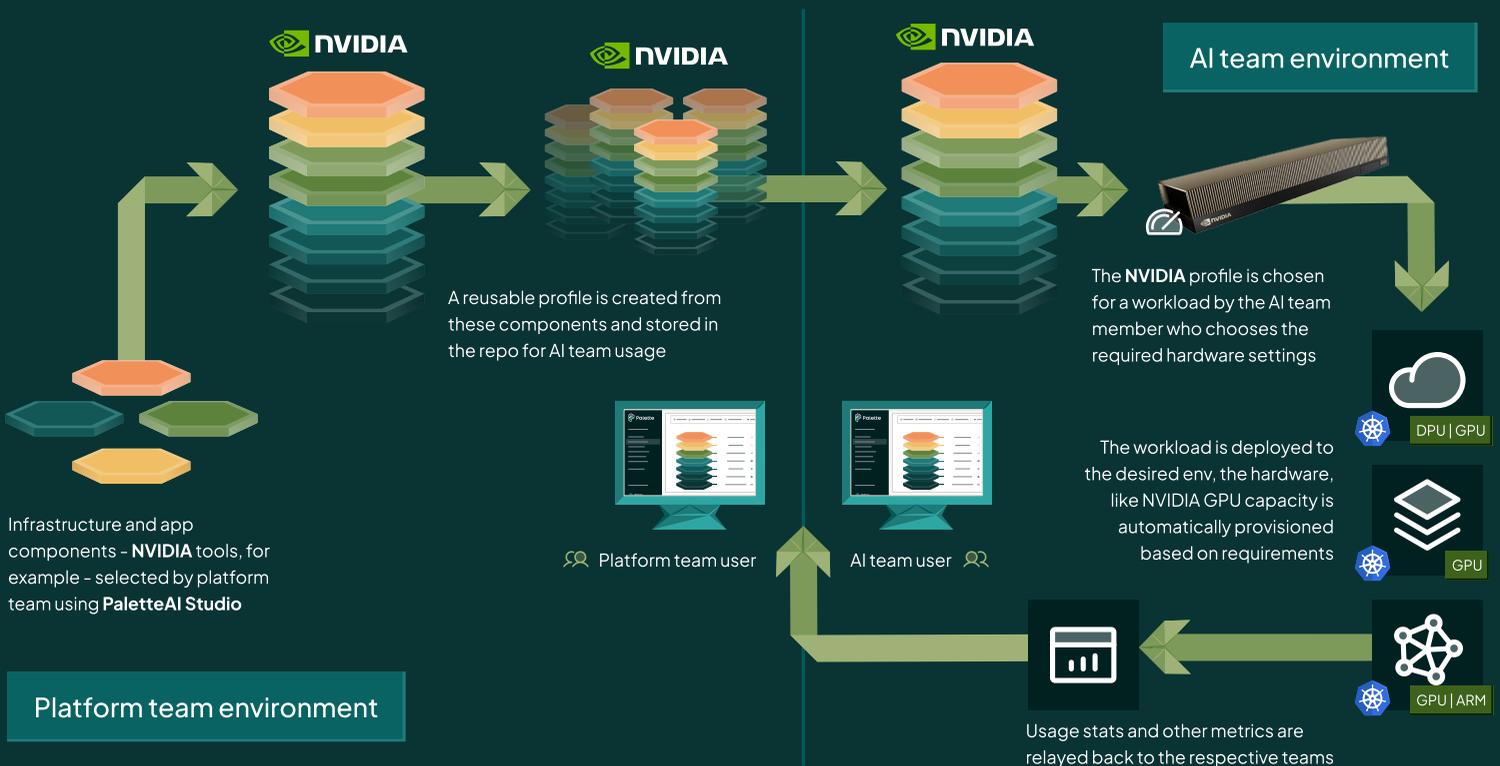
Platform teams also curate workload profiles: preconfigured AI and ML applications enhanced with corporate IT guardrails and security policies

validated for regulated environments. Practitioners can then easily customize and deploy these workloads on the available clusters, accelerating innovation while maintaining compliance, observability, and operational consistency.

PaletteAI VerteX also supports the deployment of secure NVIDIA-powered AI factories, enabling enterprises to bring NVIDIA's AI Factory vision to life at scale while adhering to strict data protection and zero-trust requirements.

With native integration across the NVIDIA AI Enterprise suite, including NeMo and NIM, as well as NVIDIA Run:AI, PaletteAI VerteX simplifies the rollout and management of complete AI factory environments, unifying NVIDIA Blackwell GPUs and BlueField DPU automation with standardized, repeatable stacks across data centers, clouds, and the edge.

In short, platform teams can design guardrailed, FIPS-compliant templates and policies, while AI teams can deploy freely within approved and secure boundaries.



Why choose PaletteAI VerteX?

When you choose Spectro Cloud and PaletteAI VerteX, you're selecting a best-in-class solution that delivers both speed and control enabling innovation at scale with the security, compliance, and trust that regulated industries demand. We offer:

- **Accelerated ROI:** eliminating idle GPU spend, right-sizing infrastructure, and shortening pilot-to-production timelines.
- **Freedom of choice:** deploy across clouds, edge, and data center with hardened infrastructure offerings and no vendor lock-in.
- **Operational efficiency:** standardized, automated workflows and reusable blueprints reduce manual toil and human error.
- **Enterprise governance and trust:** built-in policy enforcement, auditability, and zero-trust DPU-backed security ensure consistent compliance and verified protection at scale.
- **NVIDIA-validated architecture:** PaletteAI VerteX is built on NVIDIA AI Enterprise, supporting the latest Blackwell GPUs, Grace CPUs, BlueField DPUs, and Spectrum-X networking.
- **The fast path to NVIDIA's AI Factory:** supporting evolving AI frameworks such as NIM, NeMo, and Run:AI.
- **Strong industry partnerships:** we're an NVIDIA preferred partner, and work with leaders in the AI ecosystem including HPE, Supermicro and WEKA.
- **World-class support:** we back every PaletteAI VerteX customer with global 24x7 support across your entire software and infrastructure stack.

Trusted by customers who require the most rigorous security and compliance, Spectro Cloud solutions already power mission-critical AI environments across government, defense, and tactical deployments, as well as financial services, healthcare, energy, and oil and gas.

These organizations rely on our FIPS 140-3-compliant architecture, zero-trust access model, and policy-driven automation to meet stringent data protection and operational requirements without compromising agility.

Start your journey with PaletteAI VerteX

Visit palette-ai.com to learn more about PaletteAI VerteX and arrange a live, hands-on demo tailored to your organization and use case.

