

TURN AI PILOTS INTO ALWAYS-ON BUSINESS OUTCOMES

Scale proof-of-concept wins into
production-grade services with
NVIDIA AI Computing by HPE

AI is no longer just a collection of pilots and proofs of concept (PoCs). It is becoming an integrated fabric of always-on intelligent services that people expect to work reliably and act on their behalf. To achieve this, organizations must move beyond experiments with unmanaged prompts, models, and agents that operate outside the safety net of enterprise security and governance.

Enter NVIDIA AI Computing by HPE. This portfolio of products, which includes the turnkey HPE Private Cloud AI solution, is designed to help organizations move from pilots to reliable, production-grade AI services. Co-developed by HPE and NVIDIA, HPE Private Cloud AI provides the hardware, software, and services necessary to run enterprise AI workloads at scale, giving your business a quick, efficient, and flexible way to operate AI assistants and applications across the organization.

Table 1. Accelerate the shift from AI experiments to reliable operations with HPE Private Cloud AI

AI domains	Experiment-driven approach	HPE Private Cloud AI approach
Scope of AI initiatives	One-off PoCs in isolated teams; value is difficult to measure.	A portfolio of assistants and copilots aligned to clear business outcomes, powered by enterprise data.
Data and grounding for assistants	Custom integrations and indexes for each pilot; governance is managed on a case-by-case basis.	Security, governance, and policy-driven data access, with lineage and privacy controls enforced across the platform for assistants and applications.
Model, runtime, and deployment	Each team selects models and technology stacks; switching models or scaling an application requires re-engineering.	Curated, pre-integrated models and runtimes, plus NVIDIA AI Blueprints, enable teams to deploy new assistants and endpoints quickly.
Operations, performance, and cost	Minimal observability; no standard view of throughput, latency, or token costs across pilots.	Unified observability provides insights into performance, energy use, and costs, as well as rightsizing and autoscaling for inference tiers.
Risk, compliance, and lifecycle	Inconsistent guardrails and reviews; difficult to audit prompts, outputs, and data access.	Instead of juggling experiments, you operate secure, scalable AI services that deliver measurable outcomes.



AI's potential is massive—but only if you can advance beyond pilots. With HPE Private Cloud AI, you get a turnkey platform that keeps AI reliable and available around the clock. You go from isolated experiments to secure, scalable AI services that deliver tangible outcomes.

Learn more at

[HPE.com/ai/insights](https://hpe.com/ai/insights)



Visit [HPE.com](https://hpe.com)

Learn more at

[HPE.com/ai/insights](https://hpe.com/ai/insights)

HPE

 **NVIDIA®**

[Chat now](#)

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The NVIDIA logo is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. All third-party marks are the property of their respective owners.

a00155976ENW

HEWLETT PACKARD ENTERPRISE

hpe.com