Solution brief



Secure enterprise AI with HPE, NVIDIA, and Thales for trusted production outcomes

Protect data and scale AI with confidence



Turning AI ambition into secure, compliant outcomes

As organizations move AI workloads from experimentation into production, they face a new class of risks that can slow progress or derail innovation altogether. Sensitive data used for training and inference must be protected across increasingly complex environments. At the same time, threats such as prompt injection, data exfiltration, ransomware, and model misuse continue to evolve alongside AI adoption.

Without a strong security foundation, organizations risk data breaches, compliance failures, and loss of intellectual property. This challenge becomes more acute in regulated industries where strict governance, auditability, and sovereignty requirements must be maintained at every stage of the AI lifecycle.

A clearer path to scaling AI securely

Organizations have an opportunity to move beyond fragmented approaches and adopt a solution-led model that accelerates production-ready AI while reducing risk. By integrating security directly into AI infrastructure, pipelines, and workflows, enterprises can deploy trusted AI solutions from day one rather than retrofitting controls later.

Embedding data-centric protection and runtime safeguards into AI systems enables faster deployment, improved governance, and reduced operational complexity. This creates a foundation for responsible AI innovation while maintaining control across hybrid and regulated environments.

A unified approach to secure AI from data to deployment

HPE, NVIDIA, and Thales deliver a validated, security-by-design AI solution that protects data, models, and interactions across the entire lifecycle. The solution integrates HPE Private Cloud AI, a turnkey AI factory codeveloped with NVIDIA as part of the HPE AI Factory with NVIDIA portfolio, with the Thales AI Security Fabric to provide consistent protection across data, pipelines, workloads, and AI applications.

Delivered through the HPE Unleash AI Partner Ecosystem, this validated solution reduces integration complexity and provides a faster path to production-ready AI. This unified approach combines encryption, tokenization, centralized key management, secrets management, and runtime AI protection into a single framework. It helps ensure that sensitive data and intellectual property remain secure while enabling organizations to scale AI across hybrid, multicloud, and sovereign environments.

At the core of the solution is the Thales AI Security Fabric, which unifies:

— CipherTrust Data Security Platform for encryption, tokenization, and key management

— Thales Luna Hardware Security Modules for Federal Information Processing Standard (FIPS) Validated, Common Criteria Certified hardware-rooted trust

— Imperva AI Application Security for runtime protection and prompt inspection

Together, these capabilities secure AI interactions, enforce policy controls, and reduce exposure across the full AI lifecycle.

Business outcomes delivered faster

This solution is designed to help organizations accelerate AI innovation while reducing risk and operational overhead.

Organizations benefit from faster deployments because security controls are embedded directly into AI pipelines instead of added post-deployment. This reduces manual effort and simplifies operations through a unified control plane for key management and policy enforcement.

By consolidating multiple security tools into a single platform, enterprises can help reduce operational complexity and exposure to costly breaches or compliance penalties. Consistent enforcement of encryption and access controls across environments helps improve auditability and regulatory alignment.

The solution also helps organizations:

— Protect sensitive data and intellectual property across the AI lifecycle

— Enable compliant AI innovation in regulated environments

— Maintain centralized governance with full auditability

— Help improve ROI by accelerating time to production when supported by approved proof points

Prepared for enterprise AI growth

The combined solution supports enterprise-scale AI deployments across private cloud, hybrid, and multicloud environments. It is designed to protect both structured and unstructured data, including training datasets, logs, and application data used in AI pipelines and RAG workflows.

Core capabilities include centralized data protection, runtime security, and scalable key management. The solution integrates with APIs, cloud environments, databases, and AI pipelines designed to help maintain high-throughput performance without introducing significant latency.

Runtime protections continuously monitor AI interactions, helping detect unsafe inputs, enforce policies, and help reduce the risk of data leakage or misuse. This ensures that AI systems remain both secure and reliable in production environments.

Key use cases across industries

This solution supports a broad range of industry and cross-functional use cases where security, compliance, and data control are critical.

In financial services, organizations can protect sensitive financial data and ensure compliance with regulatory mandates while enabling AI-driven insights. In healthcare and life sciences, the solution helps safeguard Protected Health Information (PHI) and support secure model development for clinical and research use cases.

Regulated government and public-sector organizations can deploy AI in sovereign or air-gapped environments with full control over encryption, key ownership, and policy enforcement.

Across industries, common use cases include secure AI pipelines for regulated data, centralized key management, trusted AI data repositories, and protection of AI applications from threats such as prompt injection and model misuse.

The partner ecosystem advantage

The HPE Unleash AI Partner Ecosystem brings together HPE infrastructure, NVIDIA acceleration, and trusted partner data security solutions, such as Thales, to provide preintegrated, outcome-focused AI solutions. This solution provides a validated approach that reduces the complexity of building AI environments from scratch, reducing integration risk, accelerating deployment timelines, and moving from pilot to production faster.

By aligning infrastructure, acceleration, and security into a single solution, organizations gain a clear and governed path to enterprise AI at scale.

Built on a trusted AI foundation

This solution runs on HPE infrastructure designed to support enterprise AI workloads at scale. HPE Private Cloud AI and HPE AI Factory with NVIDIA provide the underlying compute, storage, and orchestration capabilities needed to run AI applications across hybrid and private environments.

Thales integrates directly into this architecture, delivering centralized security controls across data, pipelines, and workloads. Encryption, tokenization, and key management operate consistently across environments while hardware security modules ensure a strong root of trust.

NVIDIA provides accelerated computing capabilities that enable high-performance AI training and inference. Together, HPE, NVIDIA, and Thales deliver a balanced architecture that combines performance, security, and scalability.

Secure AI with integrated protection

Deploy production-ready AI with built-in security, governance, and performance from day one

Ready for trusted AI at scale

Organizations choose this solution to accelerate AI innovation without compromising security or compliance. By embedding protection across the entire lifecycle, enterprises can reduce risk while improving speed to value.

The combination of HPE, NVIDIA, and Thales delivers:

— A unified, security-first approach to enterprise AI

— Faster deployment with integrated infrastructure and security

— Reduced operational complexity through centralized control

— Scalable performance across hybrid and regulated environments

This integrated approach allows organizations to confidently move AI workloads into production while maintaining governance, compliance, and trust.

Thales AI Security use case for HPE Private Cloud AI and HPE AI Factory

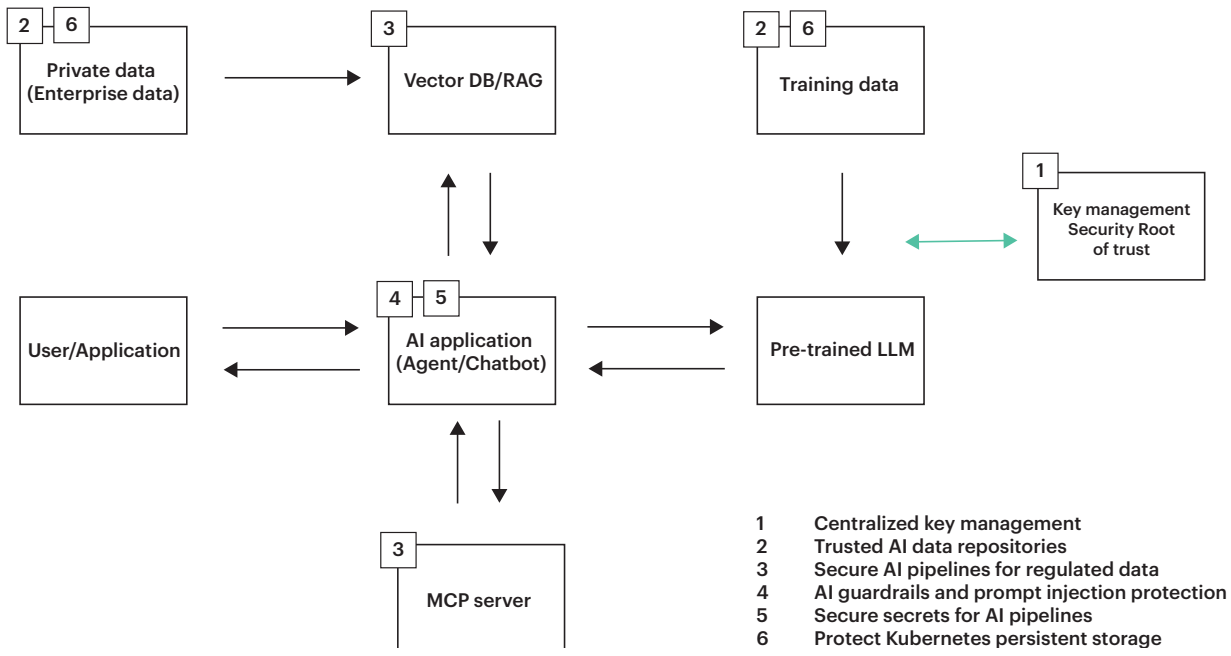
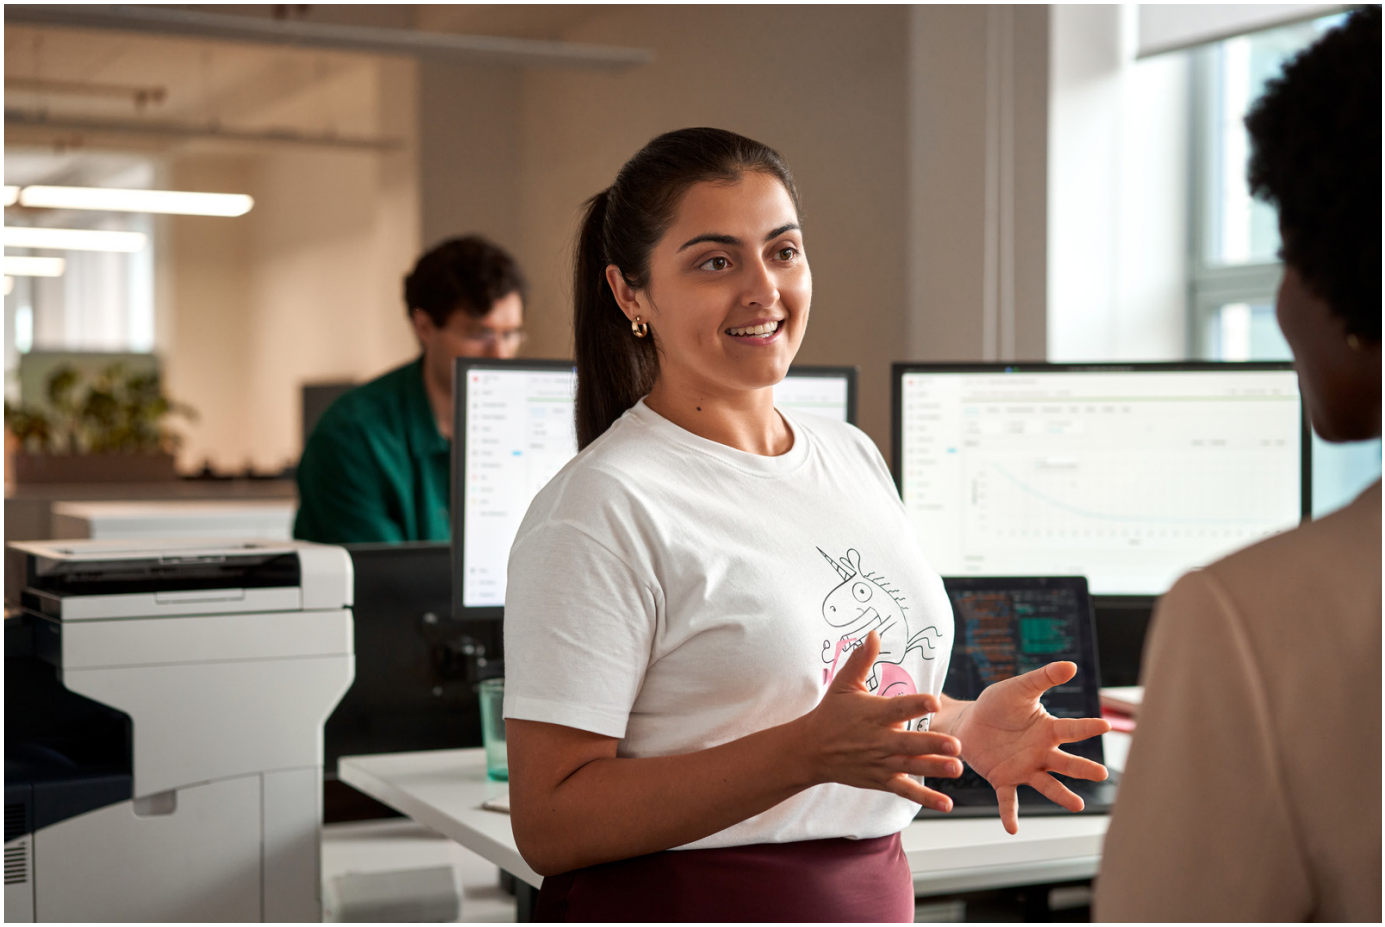


Figure 1. A data-centric security approach helps organizations maintain continuous protection, governance, and control of sensitive information across the AI environments, regardless of changes to models, tools, or architectures. Image courtesy of Thales.



Visit HPE.com

Learn more at

HPE Unleash AI Partner Ecosystem

HPE.com/Private-Cloud-AI

HPE AI Factory with 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.

a00162680ENW

HEWLETT PACKARD ENTERPRISE

hpe.com