Deploy AI securely across organizational boundaries

What's the offer?

To remain competitive, organizations need to use AI with sensitive data like financial information, healthcare data, intellectual property, and proprietary models. However, traditional security controls leave data exposed while it is being processed. Unlike conventional approaches that secure data only at rest and in transit, Fortanix Confidential AI protects data, prompts, intellectual property (IP), models, and results while in use, within an approved, cryptographically verified runtime.

With Fortanix Confidential AI, organizations can use sensitive data for AI on-premises, in private clouds, and in regulated, sovereign, or air-gapped environments with greater security, compliance, and trust.

Industry use cases

Target industries and use cases:

- **Education and research:** Secure AI factories and research environments for advanced AI with sensitive datasets, IP, and regulated research data
- **Federal and public sector:** Secure AI for mission workflows, intelligence analysis, multilingual transcription, citizen services, sovereign AI, and classified, air-gapped, or disconnected environments
- **Financial services:** Private AI for customer support, fraud detection, trading-floor communications, regulated-record review, risk analysis, model governance, and internal knowledge workflows
- **Healthcare and life sciences:** Clinical documentation, care-team voice interfaces, patient-service workflows, research analytics, and AI use cases involving protected health information, sensitive research data, or regulated clinical environments
- **Sovereign AI, neoclouds, and service providers:** Private or sovereign AI services that require data residency, jurisdictional control, model protection, and verifiable execution
- **Model providers, AI application providers, and ISVs:** Secure deployment of proprietary, customized, self-hosted, or domain-specific AI models and AI-enabled applications into customer-controlled environments

Business value

- Unlock high-value AI use cases that would otherwise be blocked
- Accelerate AI adoption by establishing trust in how data and models are protected
- Deliver a production-grade private AI platform aligned with enterprise, regulated, and sovereign requirements
- Simplify compliance with data privacy, residency, and governance requirements
- Enable customers, model providers, and AI application providers to collaborate securely

More security, strong performance

<10% throughput overhead

in modern confidential computing environments[1]

[1]Cornell University, September 2025, "Confidential LLM Inference: Performance and Cost Across CPU and GPU TEEs."

Fortanix Confidential AI: Technical architecture

Deploy in fully air-gapped and connected sovereign environments

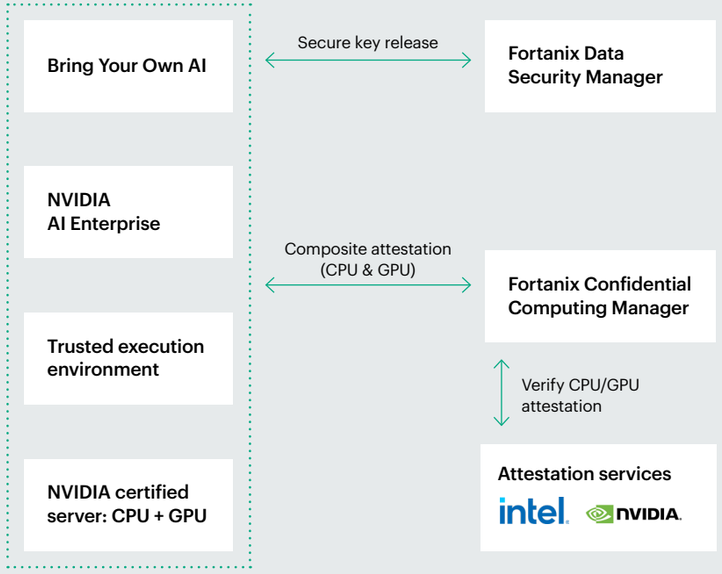

Customer pain points and solution benefits

Pain point	Solution benefit
High-value AI use cases are blocked because traditional security controls don't protect data and apps at runtime	Run AI models, applications, agents, and sensitive data in a Confidential Computing environment, where protected workloads unlock only when runtime, identity, policy, and key-release conditions are verified
Model providers, AI application providers, and ISVs are reluctant to deploy proprietary models or AI-enabled applications into customer-controlled environments	A policy-controlled, cryptographically enforced execution boundary lets model IP, application logic, and customer data be used together without exposing the model or application provider's assets to the customer environment or the customer's sensitive data to the provider
Regulated, sovereign, and public sector customers need AI deployments that preserve data residency, jurisdictional control, and operational sovereignty	The solution supports private, on-premises, sovereign, and air-gapped AI deployments with external key custody, cryptographic policy enforcement, and audit evidence
Compliance and governance teams need proof that protected AI workloads ran only in approved environments	Audit evidence from attestation, policy decisions, and key-release events give teams verifiable records of what ran, where, under what conditions, and who or what accessed protected assets
Enterprises struggle to move AI from pilot to production because security, infrastructure, and governance controls are not integrated	The solution delivers a deployable production architecture with integrated runtime security, governance, and control
Customers need to support multiple AI deployment patterns, including proprietary or customized models, self-hosted applications, and domain-specific workloads	A consistent runtime trust model across AI models and applications helps customers secure different workload types, deployment models, and operating environments without redesigning governance for each use case
Implementing new solutions can bring vendor integration challenges	Solution-in-a-box: Fully validated, integrated (hardware and software) solutions built at an HPE factory. Right-sized configurations to fit any deployment
Unpredictable pricing models and the need to manage time-consuming, complex procurement processes	Ready-to-quote same day with competitive prices

Solution components

- Fortanix Confidential AI establishes a policy-controlled, cryptographically enforced execution boundary for AI workloads, anchored by two components.
  - Fortanix Data Security Manager (DSM) holds the encryption keys protecting models and data, safeguarded in FIPS 140-2 Level 3 hardware. Keys never leave custody until the conditions for Secure Key Release are met.
  - Fortanix Confidential Computing Manager (CCM) determines when those conditions are met. It performs composite attestation of both CPU and GPU, evaluates the evidence against customer-defined policy, and records audit evidence for every decision, protecting data and applications while in use, not only at rest and in transit.
- HPE Private Cloud AI provides the private, production-ready infrastructure foundation for deploying confidential AI workloads in customer-controlled, regulated, sovereign, and air-gapped environments.
- Model providers and AI application partners provide the AI model or application layer.

More resources

Fortanix Confidential AI
Building Confidential Inference Systems: Securely Deploy Frontier Models On-Prem and Protect Enterprise Data
Fortanix Introduces Confidential AI to Advance the Development of Richer AI Models and Applications press release
HPE Private Cloud AI

Visit HPE.com

HEWLETT PACKARD ENTERPRISE
hpe.com

a00161061enw