Solution brief

Enabling sovereign AI initiatives

Operationalizing AI with security, compliance, and control

AI is reshaping digital innovation and operational strategy across governments and regulated industries. Scaling AI responsibly requires platforms that combine the speed and flexibility of cloud with built-in support for data sovereignty, intellectual property and data protection, and governmental and organizational regulatory compliance.

Designed for governments, regulated industries, and enterprises with strict sovereignty requirements, HPE provides a secure, compliant, and customer-controlled environment for building and operating AI at scale. As part of the NVIDIA Computing by HPE portfolio, HPE AI Factory sovereign co-engineered with NVIDIA, combines HPE hardware, software, an intelligent control plane, and services with the latest NVIDIA accelerated computing, software, and networking. As a result, organizations can use this platform to deploy AI across its entire lifecycle while ensuring data, models, and operations adhere jurisdictional and organizational control.

Understanding control in the AI ecosystem

Data sovereignty, AI sovereignty, and sovereign AI are related concepts, but they address different layers of control in the AI stack.

- **Data sovereignty** focuses on where data resides and which laws govern it. It ensures that sensitive data—such as personal, healthcare, financial, or government information—is stored, processed, and accessed in accordance with the legal and regulatory requirements of a specific jurisdiction. The primary concern is compliance, privacy, and protection from unauthorized or extraterritorial access.
- **AI sovereignty** extends beyond data to include control over the AI lifecycle itself. This includes who owns and governs the models, algorithms, training processes, and decision logic, as well as how AI systems are developed, deployed, and updated. AI sovereignty ensures that organizations or nations retain authority over how AI behaves, how outcomes are generated, and how intellectual property and policy objectives are upheld—rather than relying on opaque or externally controlled systems.
- **Sovereign AI** is the practical implementation of these principles. It combines data sovereignty and AI sovereignty into an operational model that delivers AI capabilities within a trusted, jurisdictionally controlled environment. Sovereign AI enables organizations to build, train, deploy, and operate AI with full control over data, models, infrastructure, and governance—often providing a cloud-like experience while meeting regulatory, security, and national or industry-specific requirements.

The sovereign AI strategic imperative

Sovereign AI addresses the most critical data risks organizations face as AI moves into production—ensuring control over who can access sensitive data, where it resides, and how it is used. According to Stanford's "2025 AI Index Report"[1], AI incidents jumped by 56.4% in a single year, underscoring the urgent need for robust safeguards.

By maintaining jurisdictional control, sovereign AI helps prevent extraterritorial access to data, reduces the risk of data leakage through shared or multi-tenant AI services, and supports compliance with strict data residency and privacy regulations such as the Health Insurance Portability and Accountability Act (HIPAA). At the same time, it protects against technical AI threats, including data poisoning during model training and model inversion attacks that attempt to extract sensitive information from AI outputs—delivering a secure, cloud-like AI experience built for regulated and mission-critical environments.

[1] "2025 AI Index Report", Stanford, 2025.

Overview: HPE AI Factory sovereign, co-engineered with NVIDIA

HPE AI Factory sovereign brings sovereign AI to life by providing a purpose-built, enterprise-grade platform designed to operate within trusted, jurisdictionally controlled environments. As a result, it enables organizations to build, train, deploy, and manage AI workloads with full control over data, models, infrastructure, and governance. By integrating security, compliance, and lifecycle management into a single AI factory model, HPE empowers governments and regulated industries to operationalize AI with confidence, turning sovereignty from a policy requirement into a practical, scalable reality.

The backbone of HPE AI Factory sovereign is the infrastructure on which the AI workloads are developed and deployed. HPE AI Factory sovereign is designed to solve an organizations' specific challenges with a fully integrated stack including hardware, software, control plane, storage, networking, and services. The sovereign-ready architecture features compliance layers, isolated environments, and a secure manufacturing pedigree.

HPE AI Factory sovereign solutions are built using NVIDIA-powered AI infrastructure as a foundational component and leverage NVIDIA accelerated computing, NVIDIA® Blackwell Ultra GPUs (HGX B300), NVIDIA networking technologies, and the NVIDIA AI Enterprise software suite. This also includes NVIDIA AI Factory for Government which provides a pre-validated, full-stack reference design that solves the complexity of building sovereign, high-performance AI infrastructure while meeting rigorous government security and compliance standards like FedRAMP and NIST.

Deployed on-premises, HPE AI Factory sovereign integrates with existing enterprise or government infrastructures. In addition, HPE AI Factory sovereign provides hard multi-tenancy and compliance-driven architecture that ensures the highest level of security, control, and ownership of sensitive data across the entire AI lifecycle, while mitigating the risk of regulatory non-compliance.

Built on three foundational pillars—security, compliance, and control—HPE AI Factory sovereign delivers trusted AI infrastructure without sacrificing performance, scalability, or innovation.

Three pillars for a sovereign AI factory

1. Security

HPE AI Factory sovereign gives customers full control over their sensitive data and infrastructure, keeping it within borders:

- **Consulting services:** HPE offers a full complement of services to help customers integrate sovereign security controls into their frameworks and ensure alignment with complex local regulations and compliance standards. Capabilities span security, governance and risk assessment, compliance and regulatory advisory, security strategy and roadmapping, vulnerability and threat insights, and digital trust and resilience support.
- **On-premises data control:** HPE AI Factory sovereign processes and stores data locally, so sensitive data, including critical national or business security information, remains within a country's or organization's borders and is subject to local laws.
- **Air-gapped:** For highly regulated industries or national security applications, HPE AI Factory sovereign offers customers completely isolated critical infrastructure and data from external networks.

With a trusted sovereign AI environment for sensitive workloads in place, governments and regulated organizations can reduce the risk that data, intellectual property, and critical infrastructure systems are exposed. This strict security protocol greatly reduces the risks of data breaches and cyberattacks on IP systems and protects against leaks of classified national information.

2. Compliance

HPE AI Factory sovereign is explicitly designed to meet compliance regulations and stringent certifications such as FIPS, NIST, CMMC, FedRAMP, and local standards like SecNumCloud while also providing product sovereignty and solution-level compliance support.

- **Services for sovereign AI:** HPE offers cybersecurity and advisory services to assist organizations with AI governance, risk management and compliance with General Data Protection Regulation (GDPR) and HIPAA regulations. These services help integrate sovereign security solutions into a risk framework for regulatory alignment. In addition, HPE's advisory, security, risk, and compliance services help organizations understand and address gaps in their security posture, governance, and compliance programs while HPE Managed Security Services provide continuous monitoring, detection, and response capabilities, helping organizations strengthen protection across infrastructure and hybrid environments.
- **Localized infrastructure:** HPC/AI infrastructure is deployed on-premises or in a data center in the customer's country or region. The design addresses data sovereignty laws such as GDPR, by preventing data transfers across borders.
- **Air-gapped management:** For customers with the strictest security needs, HPE offers air-gapped management capabilities, removing external network exposure.

Because a sovereign AI architecture is deployed locally, organizations can navigate the regulatory compliance of data sovereignty laws much easier and more confidently. As a result, they can approve and deploy AI initiatives with faster time to value. Additionally, the risk of regulatory non-compliance and penalties are greatly reduced since sovereign AI infrastructures operate securely behind the organization's own digital borders.

3. Control

HPE AI Factory sovereign enables real-time insights into the entire AI lifecycle, including data lineage, model behavior, policy enforcement, and system health.

- The unified control plane acts as a central layer for management, orchestration, and automation. For monitoring, it uses "checks"—individual monitors created for every virtual machine or container—to track health and performance metrics.
- The unified control plane also provides full-stack observability for HPE AI Factory sovereign. It is designed to use AI to observe AI, giving deep visibility across the infrastructure and the AI workloads running on it.
- HPE provides a comprehensive sovereign HPE AI Factory service for organizations, offering secure infrastructure, specialized software, expert consulting, implementation, and support designed for data, technological, and operational sovereignty within defined borders.

Customer-controlled sovereign AI infrastructure and data environments give unprecedented full visibility into AI workloads and usage. As such, management can ensure that all users and system stakeholders are compliant with data and network privacy, as well as detect and respond immediately to unauthorized network access. This sovereign approach to centralized management across hybrid environments, policy-based access, and workload control also gives organizations the flexibility to quickly adapt their AI environment to changes in national or business policies and regulatory laws.

Architecture and key capabilities

Unlike standard AI platforms, which often rely on logical separation and shared hardware, HPE AI Factory sovereign's hard multi-tenancy means that each customer or department gets physically separated compute, storage, and network resources. This avoids risks from "noisy" or "nosy" neighbors and ensures that sensitive workloads are truly isolated. This distinction is vital for regulatory compliance and for protecting against insider threats or cross-tenant vulnerabilities.

HPE AI Factory sovereign, co-engineered with NVIDIA solution stack

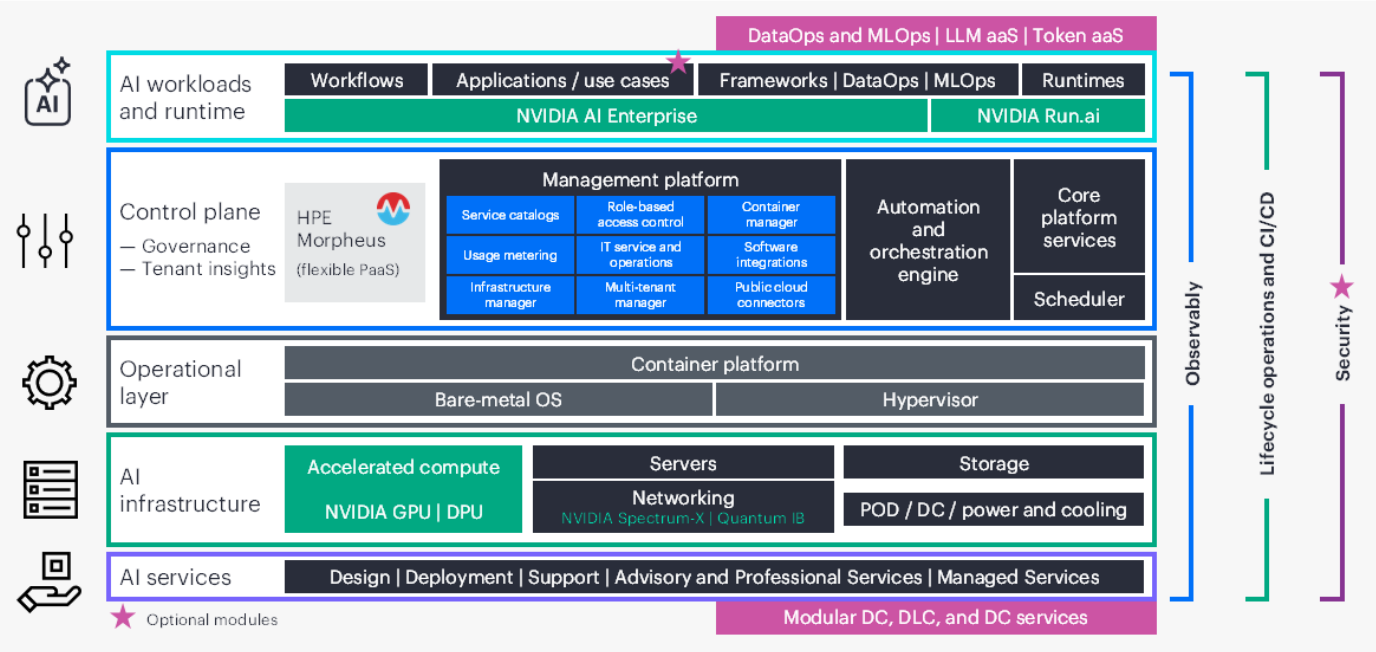

HPE AI Factory sovereign uses only hardware that meets national security guidelines, and the supply chain is tightly controlled. This reduces the risk of hardware-based vulnerabilities or tampering, a growing concern as critical infrastructure increasingly relies on foreign or third-party components. For governments and regulated industries, this level of assurance helps mitigate national security risks.

The system is designed to meet stringent certifications such as FIPS, NIST, CMMC, FedRAMP, and local standards like SecNumCloud. Compliance is enforced not just at the product level, but across the entire solution—including integration, operational practices, and ongoing monitoring. This is crucial for organizations subject to audits, regulatory oversight, or classified data requirements, ensuring they remain in good standing and avoid hefty penalties.

Real-world applications

Besides meeting sovereign AI challenges for government agencies and public sectors, HPE AI Factory sovereign can be customized for industry-specific compliance. The sovereign AI architecture supports full lifecycle AI operations, including model training, computer vision, and GenAI services, with compliance and data security features necessary for regulated environments. For example, HPE AI Factory sovereign can deliver security, compliance, and control for classified research in defense, or privacy-sensitive workloads in healthcare, and other data-sensitive use cases:

Government and regulated industries

- High-security AI systems for data residency compliance
- Strictly controlled data movement for smart nation programs
- Secure AI systems for National Resilience initiatives (emergency services, public safety, border control)

Financial Services

- Complete and secure observability for auditing and regulating AI model development
- Strict compliance systems for fraud detection, risk scoring, and other data-sensitive customer services

Healthcare and Life Sciences

- Highest-level security and compliance for sensitive clinical data in medical AI system development, including biobank analysis.

Take the next step

Ready to get started? Explore purchasing options or engage with HPE experts to determine the best solution for your business needs.

Learn more at

hpe.com/us/en/ai-factory

hpe.com/ai

Visit HPE.com

Achieving sovereign AI across your organization

Beyond technology, HPE brings deep expertise and comprehensive services to every engagement. HPE Services teams partner closely with organizations from initial strategy and planning through deployment and ongoing operations. Whether supporting solution design, financing, education, or long-term optimization, HPE helps accelerate implementation and ensures AI factories can evolve as business needs change.

Start your sovereign AI journey with HPE Services assessing your organization's environment and processes, identifying gaps, assigning ownership for controls, and the necessary compliance laws and regulations. As AI regulations and technology evolve, HPE provides continuous compliance and certification assistance, modular upgrades, and support for advanced security features such as confidential computing. Services include workshops, risk assessments, architecture analysis, and knowledge transfer to ensure organizations remain agile and compliant as regulations and technology evolve.

Relying on standard AI platforms without sovereign features heightens an organization's risk exposure to security vulnerabilities, regulatory non-compliance, and loss of control over sensitive data. As part of the NVIDIA Computing by HPE portfolio, HPE AI Factory sovereign is designed to mitigate these risks by integrating security, compliance, control, support services, and operational sovereignty at every layer, supporting national priorities and classified workloads throughout day-to-day operations.

Chat now

a00156328ENW

HEWLETT PACKARD ENTERPRISE

hpe.com