WHITE PAPER

Securing the AI Data Center: Why CIOs and CISOs Must Prepare Now

How Fortinet Secures the AI Data Center Across Network, Model, and Governance Layers

Executive Summary

AI is transforming the enterprise data center, driving unprecedented demands for scale, speed, and security. The shift from predictable north-south traffic to massive east-west data movement between GPUs, APIs, and storage clusters has created a new class of risk. Legacy architectures simply cannot provide the inspection, segmentation, or policy control needed to protect sensitive data, models, and workloads in motion.

Fortinet's Secure AI Data Center solution delivers the visibility, performance, and protection required for this new reality. Anchored by FortiAIGate and powered by Fortinet ASIC technology, it combines hyperscale firewalling, zero-trust segmentation, and AI-aware runtime inspection into a unified security fabric purpose-built for AI operations.

Fortinet's solution is grounded in the Secure AI Foundation Framework (SAIFF), which defines protection across four layers: infrastructure and data security, application and API defense, model protection and runtime monitoring, and governance, risk, and compliance. Together, these layers help CIOs and CISOs secure AI innovation, manage costs, and maintain regulatory trust, building a resilient foundation for intelligent infrastructure.

The AI Data Center Revolution

The evolution of the data center is driven by the scale and type of traffic generated by AI workloads. Unlike traditional north-south flows, AI requires vast, unpredictable east-west traffic between GPU clusters, storage systems, and inference APIs. Every model update, dataset transfer, and API query is now a high-volume transaction that legacy infrastructure struggles to inspect without creating bottlenecks.

To meet this demand, the modern AI data center must evolve its network core to handle massive throughput and ultra-low latency while protecting the sensitive data (training/context), large language models (LLMs), and continuous usage (API calls) that power AI services.

Unfortunately, security maturity has not kept pace, creating a gap that attackers and regulators are ready to exploit. Adding to this challenge, AI workloads are latency-sensitive, data-intensive, and exposed to new classes of threats such as prompt injection, jailbreaking, data leakage, and model poisoning or theft.

The Security Gap No One Planned For

The emergence of generative AI and LLMs has expanded the application attack surface and introduced largely unaddressed AI-related risk. This shift has created critical security gaps that legacy tools can't close.

Organizations rushing to secure their AI initiatives often end up adding a patchwork of point products, leading to a fragmented security posture. Fragmentation reduces visibility, complicates governance, and slows response times and worsens the problem of solution sprawl across AI observation, DLP, and network detection products.

The clear path forward is consolidation. Simplifying the security stack through unified platforms strengthens risk posture, streamlines operations, and lowers cost. As Gartner notes, "The AI race is driving vendor fragmentation in the short term, but will cause consolidation in the mid to long term."[1]

The challenge is that point solutions built solely for LLMs cannot provide the end-to-end protection and visibility required for today's AI-powered data centers.

The Technical Reality: Securing the AI Workload

Traditional data centers were built for predictable, structured applications and stable data flows. In contrast, AI data centers process vast, unpredictable east-west traffic between GPUs, APIs, and storage clusters. Every model update, dataset transfer, and inference query introduces new, uninspected pathways for exploitation.

Legacy firewalls, designed primarily for north-south inspection, lack the capacity to monitor high-throughput, low-latency east-west flows and cannot govern AI content (prompts or outputs).

Key threats:

- **Prompt injection and jailbreaks:** Manipulative inputs that bypass controls or trigger unintended behavior
- **Model poisoning:** Corrupted training data that degrades accuracy or embeds backdoors
- **Data leakage:** Model outputs revealing confidential or regulated information
- **Unauthorized access:** Adversaries gaining privileges to AI models or APIs
- **Excessive consumption:** Resource-intensive queries that degrade performance or inflate costs

Strategic Shift: The Secure AI Foundation

To address these challenges, CIOs and CISOs have begun to align on three security priorities:

- Governance of AI adoption and risk
- Cost management
- Model and data security

This has resulted in the **Secure AI Foundation Framework**, a defense pyramid spanning infrastructure and data security at its base, application and API protection (WAF, prompt-injection defenses) at the next tier, model protection and runtime monitoring, and governance, risk, and compliance at the top.

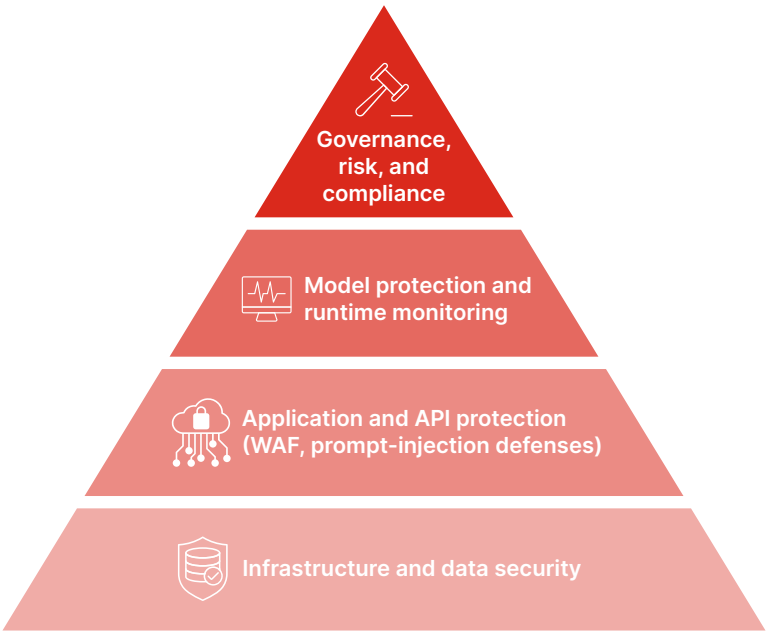

Figure 1: The Secure AI Foundation Framework (SAIFF)

Secure the Model, the Data, and the Workflow

AI data centers demand security aligned with how AI operates at three concurrent layers:

- **Perimeter and segmentation:** Control every entry and exit point. GPU clusters and inference nodes must be isolated and continuously verified with granular identity and privilege controls.
- **Data and application protection:** APIs must be shielded from adversarial queries, theft, and unauthorized access.
- **AI runtime security:** Continuous inspection of model inputs and outputs prevents malicious prompts, data leakage, and compliance breaches.

Introducing the Secure AI Data Center Solution

Fortinet has engineered a new class of protection based on these priorities: the **Secure AI Data Center solution**. Anchored by **FortiAIGate**, it protects the network, data, and model by combining hyperscale firewall performance, AI runtime inspection, and zero-trust segmentation, all purpose-built for AI workloads.

Core capabilities:

- **AI runtime security:** Inline inspection of prompts and responses to detect injection, exfiltration, and unsafe content
- **Zero-trust segmentation:** Hardware-accelerated isolation between GPU clusters and data zones
- **Granular identity and access control:** Privilege-based model access
- **Performance and sustainability:** ASIC-powered throughput with reduced energy cost

Why CIOs and CISOs Need to Act Now

CIOs view AI as a growth catalyst; CISOs see an expanding attack surface. Both perspectives are valid. However, unprotected AI models are rapidly becoming a leading enterprise risk. At the same time, regulations such as the EU AI Act and the NIST AI Risk Management Framework (AI RMF) are raising the bar for control and auditability.

Delaying AI-specific protection isn't just risky. It's negligent. Forward-looking leaders are looking to embed zero-trust segmentation, deploy runtime inspection, and adopt ASIC-powered architectures to balance security with efficiency. Those who act now will define the trust standards competitors must follow.

From Point Products to Unified Protection

Point products, even AI "firewalls," can't effectively protect an AI data center. The resulting fragmented visibility leads to blind spots and inconsistent enforcement. Worse, AI "firewalls" only address a slice of the problem.

Competing platforms may offer broad portfolios but often lack the deep, native integration required to handle AI-level performance. Only Fortinet's Secure AI Data Center solution integrates network, application, data, and AI security into a single, high-performance, ASIC-powered fabric, delivering unified visibility, consistent enforcement, and centralized management that scales without latency penalties.

Simplifying the stack through an integrated security fabric enables faster response times, lowers operational burden, and strengthens resilience. In high-speed AI environments, where every millisecond matters, such simplification becomes a critical competitive advantage.

Fortinet's Secure AI Data Center Blueprint

AI workloads demand a complex mix of identity control, application security, LLM gateway functionality, and DLP. However, when deployed as separate tools, these isolated technologies create operational overload.

Fortinet's unified **Security Fabric** integrates network, application, data, and AI security under a single framework, solving large-scale AI challenges without management complexity.

Three pillars of protection:

1. **High-performance firewalling and segmentation**
   - ASIC-powered efficiency: Lower power per Gb than competing platforms
   - Hyperscale port density: 400 GbE interfaces for ultra-low-latency inspection
   - Deep internal segmentation: Prevents lateral movement across AI workloads

2. Integrated AI runtime security: FortiAIGate

- FortiAIGate: Protects the model itself through L7 proxy inspection and LLM-specific defenses
- FortiAIFlow: Optimizes performance and cost with semantic caching and rate-limiting
- FortiAIGuard: Inspects all LLM inputs and outputs against OWASP Top 10 risks, blocking prompt injection, jailbreaks, and data leakage

3. Intelligent LLM delivery

- Secure proxying and JWT-based authorization: Ensure legitimate access while balancing loads across back-end LLMs

Compliance and Governance Alignment

Regulations such as the **EU AI Act, NIST AI RMF**, and similar regional policies emphasize transparency, accountability, and secure model operations. The **SAIFF** aligns with these requirements through:

- **Data governance:** Applying consistent identity, DLP, and access controls across model pipelines
- **Risk management:** Enabling runtime visibility into model behavior and content safety
- **Auditability:** Centralized policy enforcement and logging through the Fortinet Security Fabric
- **Sustainability:** Power-efficient ASIC design supporting environmental compliance initiatives

This mapping helps organizations meet emerging AI assurance and certification standards while maintaining operational performance.

The Bottom Line

AI-powered data centers are quickly becoming the digital core of the modern enterprise. Without security designed for AI, however, they are also the most exposed.

Fortinet's Secure AI Data Center solution represents a fundamental shift from perimeter defense to model-centric protection, from siloed tools to unified fabrics, and from reactive compliance to proactive trust. Its ASIC-based design inspects high-throughput east-west traffic at the speed and efficiency required for AI scalability and sustainability.

CIOs and CISOs who act now with an integrated, fabric-based approach will secure their AI investments, meet evolving compliance mandates, and position their organizations to scale confidently into the next phase of intelligent infrastructure.

Next Steps

Learn more about the Fortinet Secure AI Data Center Blueprint or contact your Fortinet representative to evaluate your AI readiness.

[1] Gartner, 8 May 2025 – ID G00827855, By Radu Miclaus, Jim Hare, Leinar Ramos, Eric Goodness

November 6, 2025 11:47 AM / MKTG-2358-0-0-EN