ModelArmor

Cloud Security for AI/ML Workloads

An Open-Source Sandbox for Securing Untrusted PyTorch, TensorFlow, NVIDIA, and JupyterHub Workloads powered by KubeArmor BPF-LSM.

The Challenge: Securing AI/ML in the Cloud

Modern AI/ML development and deployment face critical security vulnerabilities.

Risk Level: **High Risk**

Pickle Module Vulnerability

Python's pickle module poses a significant security risk, potentially allowing arbitrary code execution.

Risk Level: **36%**

Adversarial Attacks

36% of AI systems face compromised outcomes due to adversarial data manipulation.

Risk Level: **Critical**

Exposed GPU/CUDA Resources

Unauthorized GPU toolkit access remains a top concern in high-performance computing environments.

Risk Level: **80%**

Container Breaches

80% of organizations using containers face misconfigurations that lead to vulnerabilities.

The Solution: ModelArmor

Secure isolation for AI/ML workloads with KubeArmor sandboxing

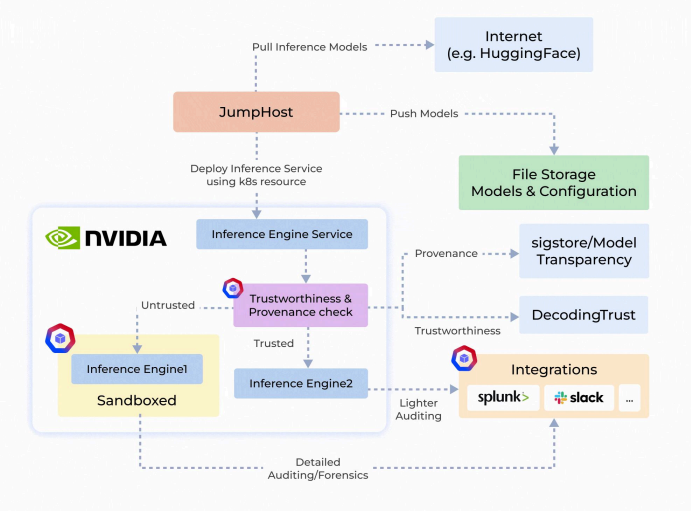

Secure TensorFlow & PyTorch

Isolated execution for TensorFlow and PyTorch models.

Container Hardening

Prevents vulnerabilities in sandboxed container environments.

Sandboxed Testing

Ensures untrusted applications execute securely.

GPU/CUDA Security

Secures NVIDIA GPU toolkits from unauthorized access.

NVIDIA GPU Protection

Prevents unauthorized access to GPU toolkits.

Framework Isolation

Separates TensorFlow and PyTorch environments.

Zero-Trust Security

Strict verification for all workloads.