White Paper

Author: John Bernard

AAEON Technology Inc.

Defense in Depth

AAEON's Multilayered Software Security Framework for Data Integrity in Decentralized Environments

April 25, 2025

Contents

Executive Summary

As artificial intelligence (AI) workloads have increasingly been moved from cloud-based environments to the edge, protecting application-layer data such as AI models, inference data, and proprietary algorithms has become a necessity, but one that has given rise to a number of challenges.

The decentralized nature of running AI algorithms on the edge, while maintaining connectivity with server devices for the necessary purposes of software updates, refined model deployment, and device management has expanded the area of exposure available to cyber threats. This is particularly true when considering the fact that edge devices lack the lack the physical security offered by centralized data centers. Moreover, the bilateral data transmission between edge devices and servers open up additional areas of vulnerability, namely cyberattacks targeting confidential data in transit.

As such, maintaining robust data security measures across entire application architecture is not possible through the use of traditional cloud-based security mechanisms.

Recognizing the need for a more comprehensive solution, AAEON has developed a multilayered software security framework designed to address security threats throughout each level of the AI application ecosystem. This framework includes server-side management tools, device-level security mechanisms, and secure tunnel communication protocols, all working in concert to maintain the integrity of critical edge AI application data.

This white paper provides an in-depth exploration of the tools that make up AAEON's software security framework, as well as how they address key cybersecurity challenges within the context of deploying AI functionality on the edge.

Understanding the Risks of Edge AI

Increased Exposure Through Decentralization

The move from cloud to edge has resulted in numerous tangible benefits when it comes to making AI functional and accessible to more industries. By running AI algorithms on the edge, applications are able to utilize real-time inferencing while reducing latency. However, edge devices are typically deployed in environments such as factories or integrated into broader infrastructure, which increases the risk of ownership for users due to their subsequent susceptibility to both physical tampering and digital intrusion.

Unlike centralized data centers with layers of controlled access, edge devices interact with the real world both in terms of digital communication with other hardware and physical accessibility, both of which increases the surface area for attacks.

Despite their physical separation from server-based management systems, each edge device can be seen as one node within a network of devices within a broader application ecosystem. Therefore, each device deployed in an uncontrolled setting is a potential entry point into this ecosystem.

Intellectual Property as a High-Value Target

Edge computing has seen a number of revolutionary innovations in recent years, from the introduction of GPUs capable of independently running extremely complex machine learning algorithms to the growth in available development resources.

As a result of these innovations, organizations have been able to utilize vast model training datasets to build algorithms capable of utilizing tailored machine learning models, application-specific inference engines, and proprietary datasets to streamline their business operations and increase the value of their commercial offering. A consequence of this is that the resulting models, once deployed, become valuable business assets.

A successful attack can that results in such assets being stolen or rendered inoperable can damage brand reputation, compromise customer privacy, and result in regulatory violations, particularly with respect to highly regulated sectors such as healthcare, finance, and infrastructure.

Advanced Threat Landscape

As the AI deployment landscape has evolved, so too has the modus operandi of cyberattacks. The techniques used to penetrate the barriers to application data have grown in sophistication, with firmware injection, supply chain exploits, man-in-the-middle attacks on communication channels, and machine learning model inversion all becoming increasingly common.

This shift renders single-solution, cloud-based cybersecurity measures inadequate to the needs of today, necessitating multifaceted security mechanisms that transcend traditional firewalls and antivirus software. Protecting AI on the edge requires multiple tools, both hardware and software-based, in order to reduce the risk of penetration across a broader area.

AAEON's Multilayered Security Framework

AAEON's software security framework has three core components, the purpose of each being to mitigate risk through specific tools deployed across the full application ecosystem:

- Server-Side Management Tools
- Device-Level Security Mechanisms
- Secure Tunnel Communication Protocols

This three-pronged approach works to protect edge devices against physical tampering, unauthorized data access from spoof attacks, the installation of corrupt or fraudulent OS and firmware, and prevent the interception of data in transit.

These layers are designed to work in tandem and streamline the performance, scalability, and manageability of AI applications at the edge.

Layer 1: Server-Side Management Tools

Certificate Authority (CA) Framework

AAEON provides a web-based UI to monitor and manage multiple edge systems from a single server. To do this, it implements a Certificate Authority (CA) framework to verify the identity and authenticity of both server and edge AI devices, the integrity of firmware and operating system updates, and whether devices are authorized to implement requested updates.

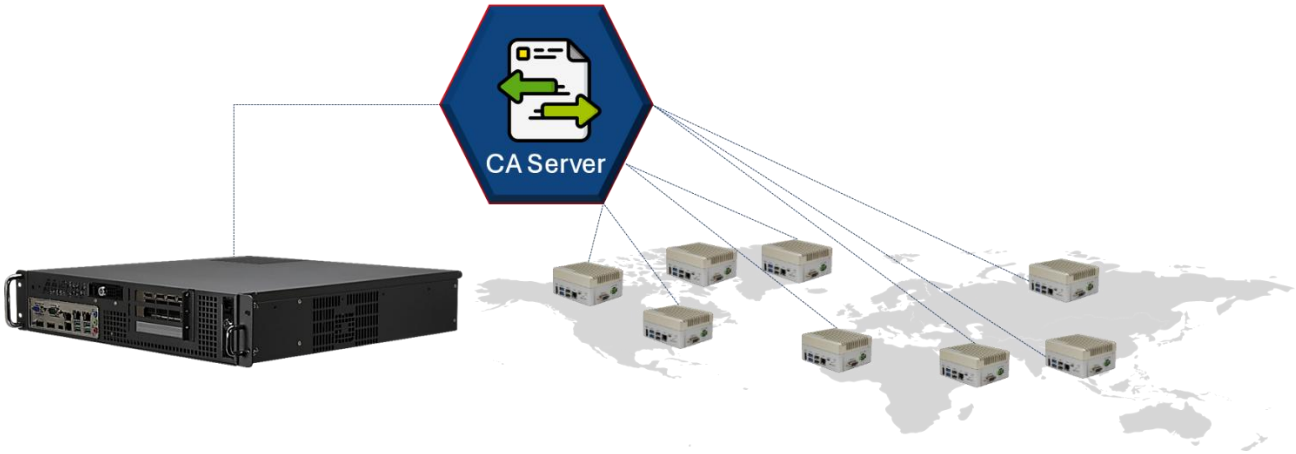

Key features include:

- **Device Authentication**: Ensures that only genuine, AAEON-made devices are able to participate in the application ecosystem, blocking access to counterfeit hardware.
- **Update Origin Validation**: Utilizes certificate chains to verify that the origin of any firmware or OS updates are trusted, guaranteeing the source of software installed is genuine.

Remote Diagnostics & Management

AAEON provides a web-based user interface for managing multiple edge AI devices, with real-time, accurate status updates to support fast diagnostics and system health monitoring.

This toolset is versatile, and provides a comprehensive platform through which users can remotely maintain edge devices, including:

- Real-time monitoring of CPU/GPU loading and frequency
- Insights into memory and storage usage
- Sensor, power, and fan diagnostics
- Group configuration of edge devices for streamlined operations

OTA (Over-the-Air) Updates

The provision of OTA updates grants users an avenue through which to securely upload new AI models, flash OS images, or perform firmware updates using the device's LAN or wireless connection to the server-side device.

The primary benefit to this is that it minimizes the need for physical intervention on the device side while providing the protection against the deployment of malicious software through CA authentication of both hardware and software prior to the update being initiated.

Key benefits include:

- Secure AI Model Deployment
- Minimal Downtime
- Malware Mitigation

Layer 2: Device-Level Security Mechanisms

Out-of-Band (OOB) Management

OOB management allows authorized users to securely recover devices in the event of system failure. By utilizing OOB management, the LAN or 4G connection to the device can be used to power the system on or off, reset the system, or utilize a Network Controller Sideband Interface (NCSI) to communicate with the system's network controller to send and receive management traffic such as updates or device health status even in the event that the primary OS is unresponsive.

Secure Boot (UEFI + BSP)

Secure Boot ensures only signed bootloaders, OS kernels, and BSP components from trusted sources can execute during system initialization, protecting the device from early-stage malware. The method for achieving this is twofold – UEFI and BSP Secure Boot sets conditions so that only authenticated firmware and software can run during system startup, while using a pre-established hardware-based root of trust to validate the digital signatures of boot components.

dTPM 2.0 Integration

A discrete Trusted Platform Module (dTPM) provides hardware-based encryption and system authentication for inter-device data transmission purposes.

File System Integration (FSI)

Combines two file systems into one packet consisting of a read-only and a read-write layer. This allows users to use two storage systems together in the form of an unchanged lower layer while storing changes, updates, or activity logs in the upper layer, thereby increasing the capacity of available storage space.

By maintaining a read-only layer, File System Integration makes it so key software such as trusted firmware or system images cannot be altered, maintaining file integrity. Conversely, in the event that the read-write layer encounters corrupted files or incorrect updates, the layer can be recreated without impacting critical files in the low layer.

Anti-Restore Protection

Automatically ensures a device is running on the most up-to-date BSP version, while also preventing hackers from reverting system software to previous versions that potentially contain known vulnerabilities that can be exploited

Disk Lock

Protects sensitive data at rest, ensuring data is inaccessible even if a device or SSD is lost, stolen, or tampered with. This is a crucial feature with respect to the prevention of physical intrusion attempts on edge devices, which can often be deployed in remote or unmanned settings.

A/B Redundancy

Utilizing an A/B partitioning scheme, A/B Redundancy allows the system to revert to a stable OS image if the primary root partition fails.

MAZU AI Model Protection

MAZU is a Trusted Execution Environment (TEE) designed to protect AI models and application data by segmenting different files, processes, and algorithms within protected execution zones.

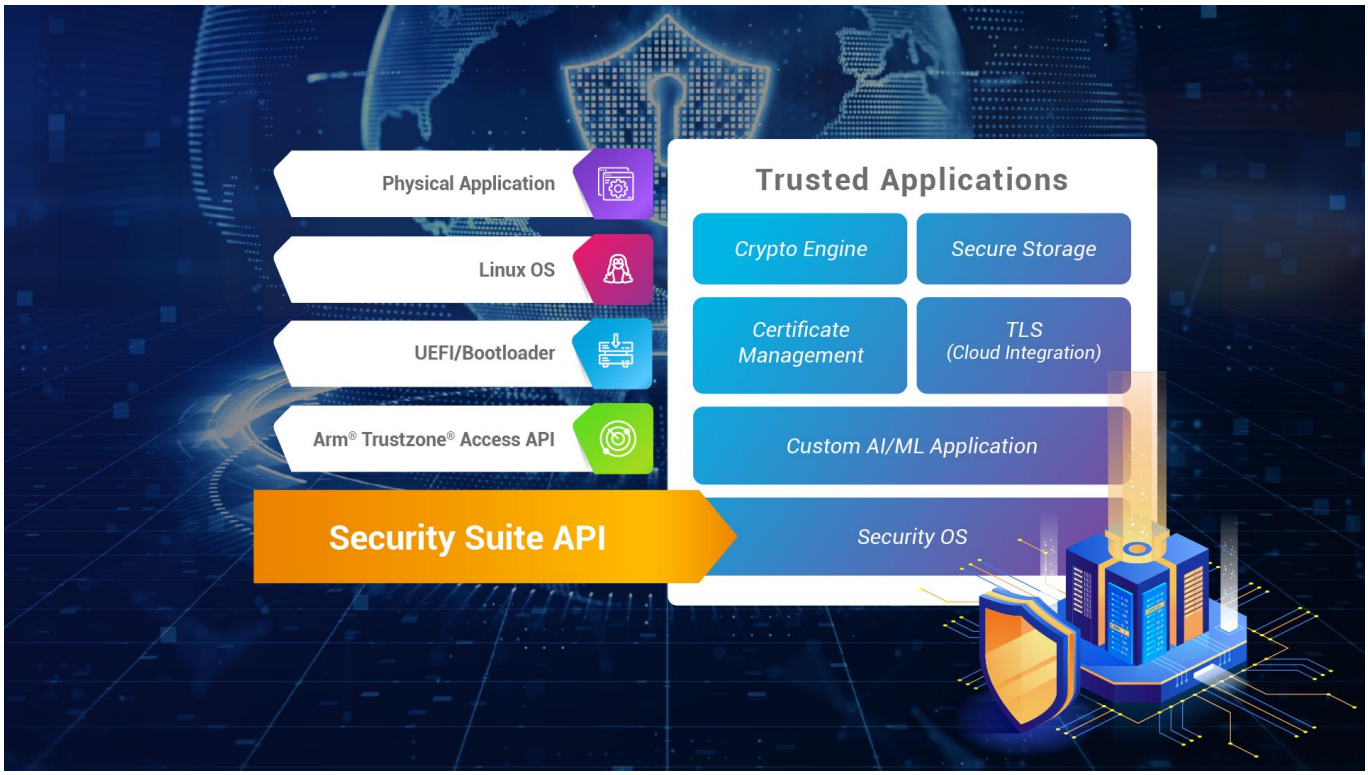

For the protection of AI models, it isolates machine learning application algorithms within a protected “Secure World”, while allowing standard applications to run in a “Normal World.” As a consequence, access to sensitive assets is gated through certified APIs with encrypted communications, secure OS environments, and certificate validation.

Key benefits include:

- **Secure Segmentation**: Machine learning models are run in isolated environments, inaccessible to the rest of the OS.
- **Controlled Access**: Only certified APIs can interact with protected AI models.
- **Secure Containers**: Provides encrypted runtime environments for inference operations.
- **Leak Prevention**: Communication channels are secured with encryption and certificate validation.

Layer 3: Secure Tunnel Communication

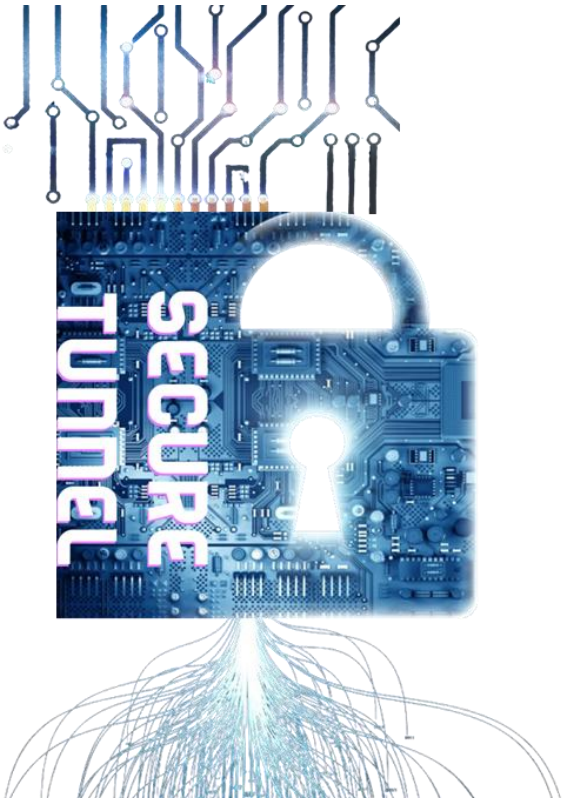

To send data between the server and edge devices, or between edge devices for software or AI model updates, AAEON uses a secure tunnel that maintains the integrity of data during transmission.

dTPM Integration for Key Exchange

By integrating a discrete TPM (dTPM) module, each device generates and stores a unique cryptographic key within secure hardware, mitigating risks of key duplication or spoofing.

Mutual Authentication

This relies on the aforementioned CA-based authentication through which the identity and authenticity of both server and edge AI devices is established as a prerequisite to data exchange being made possible.

Key benefits include:

- Man-in-the-middle attack prevention
- Sender and receiver identity verification
- Data origin integrity enforcement

Mapping Common Threats to Framework Components

The following section summarizes common security concerns for users as they relate to both edge and server environments.

Each will set out key questions or challenges faced, alongside an explanation of which specific components of AAEON's software security framework addresses them, and how.

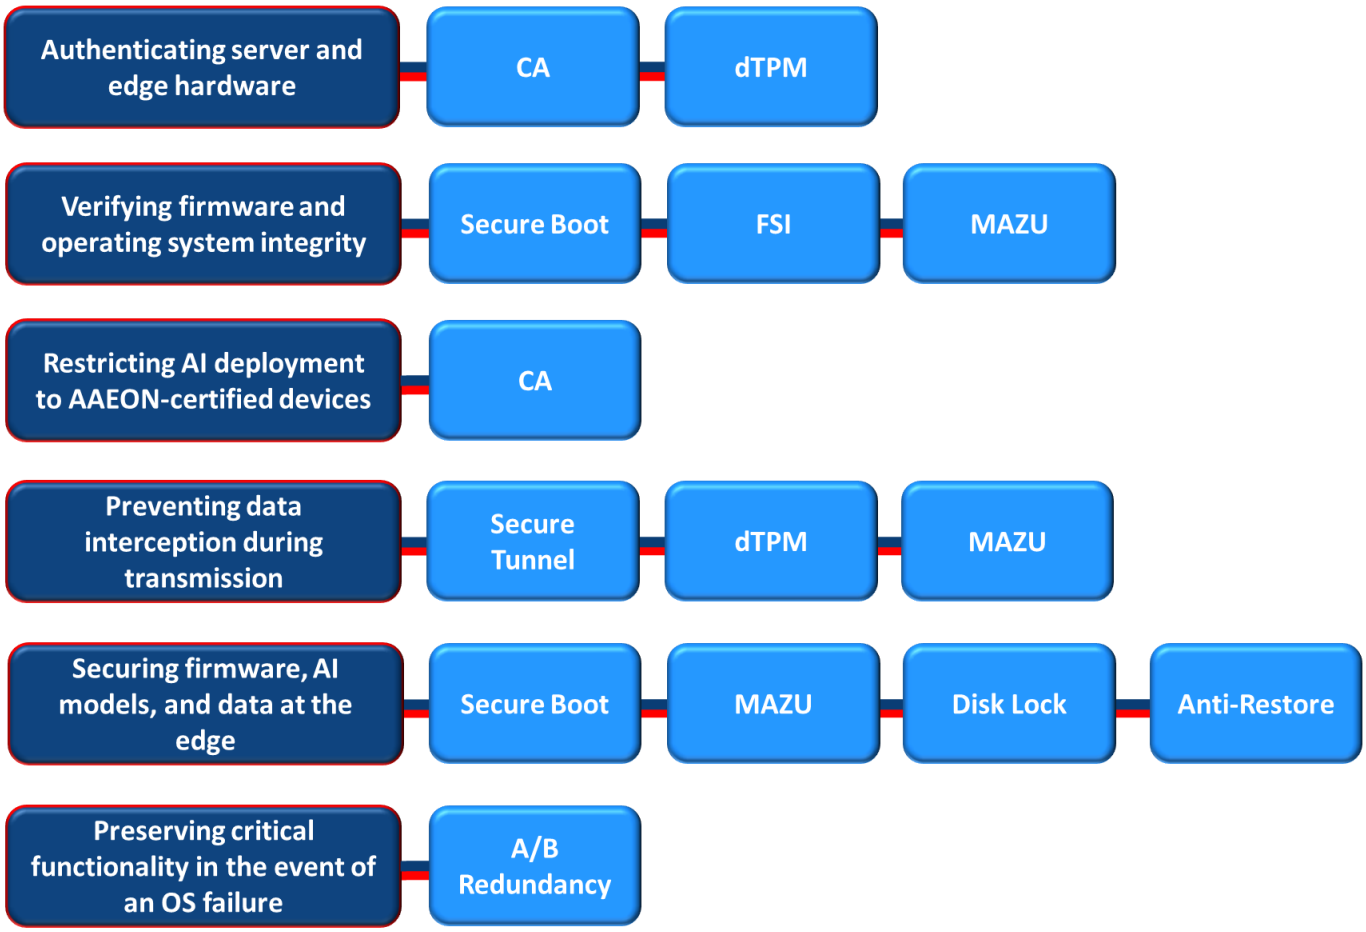

Conclusion

As decentralized edge devices become the primary environment for AI model deployment, the value and vulnerability of application-layer assets such as AI models, firmware, and data become increasingly critical. AAEON, being a leading provider of such platforms recognizes the additional challenges presented to developers in ensuring the security of such assets in modern edge computing applications.

By building a comprehensive software security framework that addresses potential vulnerabilities at both the server and edge device level of application architectures, as well as the route through which data is transmitted between them, AAEON provides a security platform that addresses the security needs of both the minutiae of segmented data in TEEs and how individual components of its security framework interact with other nodes within applications.

This framework not only protects against a wide array of cyber threats, but also preserves the integrity, authenticity, and availability of AI applications in the field. Through its robust architecture, AAEON empowers enterprises to securely deploy, manage, and scale their edge AI solutions with confidence—maintaining both operational continuity and the competitive advantage offered by proprietary AI assets.

About AAEON

Established in 1992, AAEON is one of the leading designers and manufacturers of industrial IoT and AI Edge solutions. With continual innovation as a core value, AAEON provides reliable, high-quality computing platforms including industrial motherboards and systems, rugged tablets, embedded AI Edge systems, uCPE network appliances, and LoRaWAN/WWAN solutions. AAEON also provides industry-leading experience and knowledge to provide OEM/ODM services worldwide. AAEON works closely with premier chip designers to deliver stable, reliable platforms. For an introduction to AAEON's expansive line of products and services, visit www.aaeon.com.