Comparative Study of Cybersecurity Frameworks for Ensuring Integrity of Nuclear I&C FPGA Controllers: Focusing on MITRE ATT&CK for ICS and ESTM 3.0

Jae Hwan KIM*, Kwang Seop SON, Jae Gu SONG, Se Hoon LEE, Seung Oh SEO, Young Jun LEE
Security R&D Team, Korea Atomic Energy Research Institute, 111, Daedeok-daero 989beon-gil, Yuseong-gu, Daejeon, 34057, Republic of Korea

***Keywords :** *FPGA-based Safety Controller*, Integrity Threat Modeling, MITRE ESTM 3.0, Defense-in-Depth

1. Introduction

The digitalization of nuclear power plant (NPP) instrumentation and control (I&C) systems has driven the emergence of field-programmable gate array (FPGA)-based controllers as next-generation alternatives to traditional programmable logic controllers (PLCs), offering deterministic execution, hardware-level configurability, and operating-system-free operation [1–3].

Existing frameworks, most notably MITRE ATT&CK for ICS, model adversarial behavior across operational technology (OT) networks but remain primarily anchored at the network, communication, and software execution layers, with limited coverage of hardware-level threats. The extent to which existing ICS threat frameworks can sufficiently model integrity threats targeting safety-class FPGA controllers in nuclear environments remains unclear, with direct implications for regulatory compliance and defense-in-depth strategy.

This study addresses this gap through a comparative analysis of MITRE ATT&CK for ICS and the MITRE Embedded Systems Threat Matrix (ESTM) version 3.0 (released January 2026) and makes three principal contributions: (1) a systematic mapping of structural differences between the two frameworks as applied to FPGA-based safety controllers; (2) an integrity-centric analytical perspective bridging hardware, firmware, and operational threat dimensions; and (3) a proposed dual-layer defense-in-depth strategy that extends protection to encompass firmware and hardware integrity within safety-class devices.

2. Architectural Context of FPGA Controllers

2.1 Nuclear I&C Digital Architecture

NPP I&C systems are typically structured according to the Purdue Enterprise Reference Architecture (PERA) across Levels 0 through 3 [4]. Level 0 encompasses sensors, actuators, and field devices; Level 1 hosts safety-class controllers; Level 2 hosts supervisory systems (SCADA/DCS HMIs and historian servers); and Level 3 covers plant-wide information management. Levels 0 and 1 are subject to the most stringent cybersecurity and functional integrity requirements, as they directly interface with safety-critical physical processes.

Safety-class controllers at Level 1 execute reactor trip functions and engineered safety feature (ESF) actuation deterministically, independent of higher-level system states. FPGA-based controllers are increasingly being qualified for this role, offering deterministic execution without an underlying operating system and reduced exposure to software-layer vulnerabilities, while simultaneously introducing a distinct set of hardware-level threats that necessitate a reassessment of the applicable cyber threat model [5].

2.2 Integrity-Critical Characteristics of FPGA-Based Safety Controllers

The architectural properties that qualify FPGA controllers for safety-class deployment simultaneously define four integrity-critical characteristics requiring dedicated threat consideration:

(i) OS-less execution: Logic executes directly in hardware, eliminating OS-level attack surfaces but rendering host-based monitoring and patching inapplicable — making pre-deployment integrity assurance the primary defense.

(ii) Bitstream-based configuration: All functional behavior is determined by the configuration bitstream; unauthorized modification at any lifecycle stage is undetectable by software-based controls and requires hardware-level verification mechanisms such as bitstream encryption and authentication.

(iii) Debug interface exposure: Inadequately protected interfaces (e.g., JTAG) enable unauthorized access to internal configuration and memory states, constituting a persistent hardware-layer attack vector independent of network isolation.

(iv) Supply chain dependency: Vendor-supplied synthesis tools, IP cores, and configuration utilities each represent potential Hardware Trojan insertion points; the opacity of third-party IP cores makes pre-deployment verification technically challenging.

Collectively, these characteristics define a structural threat surface operating below the abstraction layers addressed by network- and system-centric frameworks such as ATT&CK for ICS, necessitating a dedicated analytical layer extending beyond the OT network boundary.

3. Comparative Analysis of Threat Frameworks

3.1 Comparison Criteria

This study compares MITRE ATT&CK for ICS and MITRE ESTM 3.0 across four criteria spanning architectural depth, analytical philosophy, device specificity, and lifecycle scope: (1) Analytical layer — architectural depth of adversarial modeling, from control network communications to device-internal hardware interfaces; (2) Threat modeling philosophy — primarily observed, real-world TTPs versus forward-looking, theoretically feasible threat scenarios; (3) Device-level coverage — extent of coverage of embedded hardware constructs including firmware, configuration registers, debug interfaces, and hardware trust anchors; (4) Lifecycle coverage — whether the framework supports threat analysis beyond the operational phase to encompass development, supply chain verification, and security-by-design activities [6-7].

3.2 Structural Comparison

As summarized in Table 1, ATT&CK for ICS spans Purdue Model Levels 0–3 with modeling emphasis on Levels 1–3, addressing lateral movement, HMI manipulation, and process control disruption across networked assets. ESTM 3.0 targets device-level hardware and firmware interfaces analogous to Levels 0–1, enabling systematic coverage of firmware modification, side-channel attacks, and debug-port exploitation directly relevant to FPGA-based platforms.

The frameworks diverge fundamentally in philosophy and lifecycle orientation: ATT&CK for ICS is grounded in empirically documented real-world TTPs, optimized for detection engineering and operational incident response; ESTM 3.0 incorporates theoretically feasible threats alongside observed ones, supporting security-by-design and supply chain verification — a distinction critical in nuclear applications where proactive, design-phase threat modeling is essential.

At the tactic level, ESTM includes Credential Access and Exfiltration absent as standalone categories in ATT&CK for ICS — tactics that in embedded environments serve as precursors to cryptographic key recovery and malicious firmware reprogramming [4–5]. Furthermore, while ATT&CK's Inhibit Response Function addresses safety mechanism impairment at the process level, ESTM extends this to firmware and hardware interface layers, providing granular coverage of how safety function degradation manifests within a controller device.

Table 1: Key Differences Between MITRE ATT&CK for ICS and MITRE ESTM 3.0

Comparison Criteria	MITRE ATT&CK for ICS	MITRE ESTM 3.0 (Embedded Systems Threat Matrix)
Primary Target	Systems and networks (HMI, SCADA, PLC, EWS)	Components and devices (Component & Device)
Analysis Layer	Purdue Model Level 0–3 (modeling emphasis on Levels 1–3: control network and supervisory layers)	Purdue Model Level 0–1 (sensors, actuators, and controller internals)
Key Threats	Lateral movement, HMI manipulation, process control disruption	Firmware modification, side-channel attacks, debug port (JTAG) exploitation
Threat Basis	Primarily observed, real-world threats	Observed threats + proof-of-concept (PoC) and potential/emerging threats
Notable Characteristic	Optimized for detection and response during the operational phase	Optimized for Secure-by-Design and supply chain verification during the design phase

3.3 Identified Gaps in ICS-Oriented Models

The comparative analysis reveals three structural gaps in ATT&CK for ICS limiting its applicability to FPGA-based safety controller environments:

(i) Software-centric execution assumption: ATT&CK for ICS predominantly models execution techniques presupposing software-based environments (CLI, GUI, scripting engines), with limited representation of hardware-configured execution paradigms such as FPGA bitstream logic — creating a systematic blind spot at the device execution layer.

(ii) Network-oriented analytical bias: The ICS matrix primarily models attacks traversing control networks via industrial protocols, providing limited capacity to describe micro-level manipulations at configuration memory, internal buses, or hardware registers — the layers at which FPGA integrity threats materialize.

(iii) Limited firmware-layer coverage: ATT&CK for ICS does not model bare-metal environments, FPGA bitstream integrity, vendor toolchain compromise, or hardware Trojan insertion through third-party IP cores — gaps corresponding directly to the four integrity-critical characteristics of Section 2.2.

ATT&CK for ICS remains effective for network- and system-level threat modeling, but does not adequately address hardware- and firmware-centric attack vectors

against FPGA-based safety controllers, motivating the complementary application of ESTM 3.0.

4. Integrity-Oriented Threat Modeling Using ESTM 3.0

MITRE ESTM 3.0 models adversarial behavior in embedded systems by explicitly incorporating hardware-centric architectures, firmware dependencies, and device-level operational constraints — providing analytical capabilities not systematically addressed by ATT&CK for ICS, which covers these threats only incidentally. For nuclear FPGA-based safety controllers, whose security posture is fundamentally determined by firmware integrity, hardware configuration, and supply-chain trust, ESTM constitutes the appropriate primary analytical instrument at the device layer.

4.1 Threat Classes Affecting FPGA Integrity

Four primary threat classes — collectively defining the Integrity Impact Axis — converge on a single consequence: unauthorized alteration of the hardware logic governing safety-critical functions.

Bitstream Manipulation is the most direct integrity threat: adversary modification via supply-chain compromise, unauthorized physical access, or insecure update mechanisms alters controller behavior invisibly to network-level monitoring. Countermeasures: Secure Boot with hardware Root of Trust (RoT) and bitstream digital signature verification.

Debug Interface Abuse exploits development interfaces (JTAG, UART, SPI) that persist as attack surfaces in deployed systems, enabling firmware extraction, memory dumping, or full device reprogramming. Countermeasure: hardware fuse disablement or cryptographic authentication, with maintenance access under strict procedural controls.

Hardware Trojan Insertion embeds malicious HDL logic through compromised IP cores, vendor toolchains, or fabrication processes. Trojans activate under rare trigger conditions and evade functional testing, requiring side-channel analysis (SCA) for physical characterization and logic equivalence checking (LEC) for structural validation.

Side-Channel and Fault Injection Attacks remain viable in air-gapped environments: side-channel attacks extract cryptographic material via passive timing, power, or EM analysis; fault injection induces bit flips via active voltage, clock, or EM manipulation. Countermeasures: EM shielding, power filtering, and configuration readback verification or BIST circuits.

4.2 Safety Impact Perspective

Bitstream manipulation or hardware Trojan activation directly alters the FPGA logic implementing protective functions such as reactor trip and engineered safety feature actuation, causing protective actions to be delayed, suppressed, or incorrectly triggered. In multi-channel safety systems sharing a common design lineage, a single hardware or firmware vulnerability could simultaneously compromise multiple redundant channels — a failure mode functionally analogous to Common Cause Failure (CCF).

The potential for firmware-induced CCF-like cascading failure across redundant channels represents a threat that network-layer monitoring cannot detect and operational incident response cannot preempt. This reframes the nuclear cybersecurity imperative: hardware and firmware integrity assurance is not a supplementary consideration but a constitutive requirement of any complete safety defense strategy. ESTM's anticipatory, device-centric threat modeling, applied from the earliest phases of development and procurement, is therefore the appropriate analytical instrument for this threat class.

5. Integrated Defense Strategy for Nuclear FPGA Applications

No single framework provides sufficient coverage across all architectural layers relevant to nuclear FPGA-based safety controllers. This section proposes a concrete Dual-Layer defense strategy and an extension of the conventional Defense-in-Depth framework to encompass firmware and hardware integrity.

5.1 Dual-Layer Defense Architecture

The proposed architecture assigns each framework to its domain of strongest coverage, creating two mutually reinforcing layers: Layer 1 network detections serve as early warning indicators for device-level attacks, while Layer 2 integrity controls protect safety functions even when network defenses are bypassed.

Layer 1 — ATT&CK-Informed Network and System Protection addresses Purdue Levels 1–3, where adversaries exploit control network communications and industrial protocols. ATT&CK for ICS informs network segmentation, protocol-aware IDS/IPS, and access control policies, supporting detection engineering and incident response during the operational phase in alignment with current nuclear cybersecurity regulatory requirements.

Layer 2 — ESTM-Informed Device Integrity Protection addresses Purdue Levels 0–1, targeting internal hardware and firmware of embedded controllers. ESTM 3.0 informs security-by-design requirements across four control domains: (1) bitstream and firmware integrity — Secure Boot with hardware RoT and digital signature verification; (2) debug interface protection —

hardware fuse disablement or cryptographic authentication of JTAG/UART; (3) supply chain integrity — ESTM-aligned supplier requirements, SCA for physical characterization, and LEC for formal structural validation; (4) physical attack resilience — EM shielding, power filtering, and configuration readback verification or BIST circuits. Unlike Layer 1, Layer 2 must be integrated from the earliest stages of development and procurement.

5.2 Defense-in-Depth Extension Model

Conventional Defense-in-Depth strategies for nuclear I&C cybersecurity have been structured around network- and system-level controls, leaving FPGA integrity threats — which operate below the system layer — outside their scope. This study proposes extending the model to four layers, as illustrated in Fig. 1.

The upper two layers — Network and System — correspond to the existing model, addressed by ATT&CK for ICS-informed controls. The lower two layers represent the proposed extension, addressed by ESTM 3.0-informed controls: the Firmware Layer encompasses bitstream integrity verification, Secure Boot enforcement, and firmware modification protection; the Hardware Layer encompasses physical attack countermeasures, debug interface controls, hardware Trojan detection, and supply-chain assurance.

The lower two layers must be addressed during development and procurement phases, as they cannot be remediated through operational monitoring alone — making this extension a necessary evolution of cybersecurity practice for facilities deploying FPGA-based safety controllers.

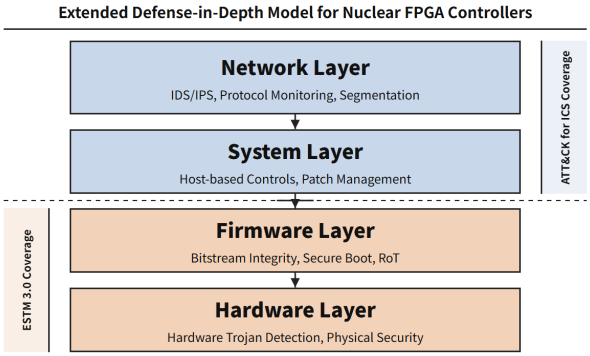

Fig. 1. Dual-Layer Defense-in-Depth Architecture Integrating ATT&CK for ICS and ESTM 3.0

6. Conclusion

ATT&CK for ICS remains necessary and effective for network- and system-level threat characterization, but its predominantly software-execution-oriented assumptions, network-oriented bias, and limited firmware coverage leave a systematic gap at precisely the layers where FPGA integrity threats materialize. ESTM 3.0 addresses this gap through structured, device-level threat modeling covering bitstream manipulation, debug interface abuse, hardware Trojan insertion, and side-channel and fault injection attacks — with an anticipatory modeling philosophy that incorporates theoretically feasible threats beyond empirically observed TTPs, well-suited to nuclear applications where proactive risk identification is essential.

The boundary of nuclear cybersecurity responsibility must be extended inward, from the control network perimeter to the silicon-level logic of safety devices themselves. The causal pathway from firmware compromise to logic alteration, safety function degradation, and potential CCF-like systemic failure demands security controls at every architectural depth. The integrated application of ATT&CK for ICS and ESTM 3.0 within the proposed Dual-Layer Defense-in-Depth architecture offers a more comprehensive analytical foundation than either framework applied in isolation, spanning development, procurement, and operational phases.

REFERENCES

[1] Idaho National Laboratory, "Cyber Risk Considerations for Nuclear Digital I&C Systems," INL/EXT-17-43201, 2022.
[2] A. Ayodeji et al., "Cyber security in the nuclear industry: A closer look at digital I&C systems," Reliability Engineering & System Safety, vol. 234, 2023.
[3] IAEA, "Computer Security Aspects of Instrumentation and Control Systems at Nuclear Facilities," IAEA Nuclear Security Series No. 33-T, Vienna, 2021.
[4] Williams, T. J. (1994). "The Purdue Enterprise Reference Architecture." Computers in Industry, 24(2-3), 141–158.
[5] U.S. NRC NUREG/CR-7006 (2010). Review Guidelines for Field-Programmable Gate Arrays in Nuclear Power Plant Safety Systems. NRC.
[6] MITRE Corp., "MITRE Launches Embedded Systems Threat Matrix (ESTM) 3.0," Jan 2026.
[7] MITRE Corp., "ATT&CK for Industrial Control Systems," 2024.