AegisSat: Securing AI-Enabled SoC FPGA Satellite Platforms

Authors: **Huimin Li**, Vusal Novruzov, Nikhilesh Singh, Lichao Wu, Mohamadreza Rostami, Ahmad-Reza Sadeghi

System Security Lab, TU Darmstadt

06/11/2025

Roadmap

01 Motivation

02 Threat Model

03 AegisSat Design

04 Implementation

05 Research Outlook

06 Conclusion

Motivation

01

Evolution of Intelligent Satellites

Satellites are shifting from static platforms to reconfigurable, intelligent systems.

Drivers: autonomy, reduced ground-control dependence, real-time decision making & bandwidth optimisation.

AI/edge computing enables real-time analytics, adaptive payload management and autonomy.

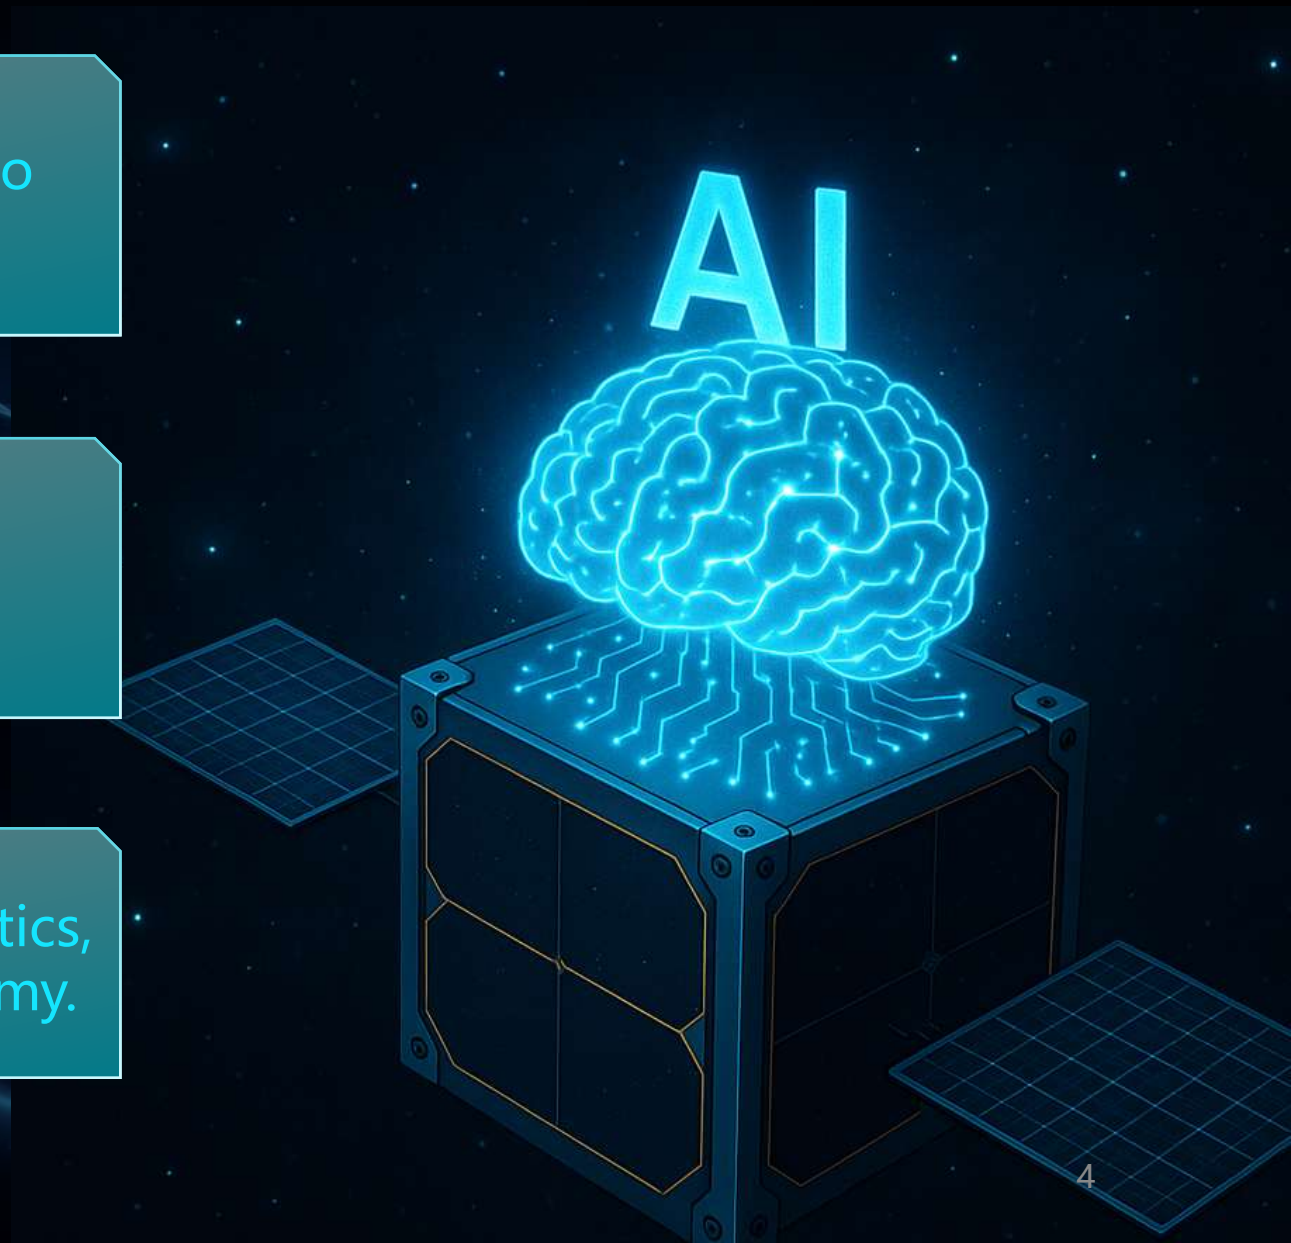

SoC FPGAs & AI Acceleration

01

FPGAs

Excel at parallelism and deterministic latency for AI workloads. Commercial off-the-shelf FPGAs provide compact size, reconfigurability and energy efficiency

02

SoC FPGAs

Integrate general-purpose processors (PS) with programmable logic (PL) in one device. Tight coupling allows complete AI processing pipelines with minimal latency.

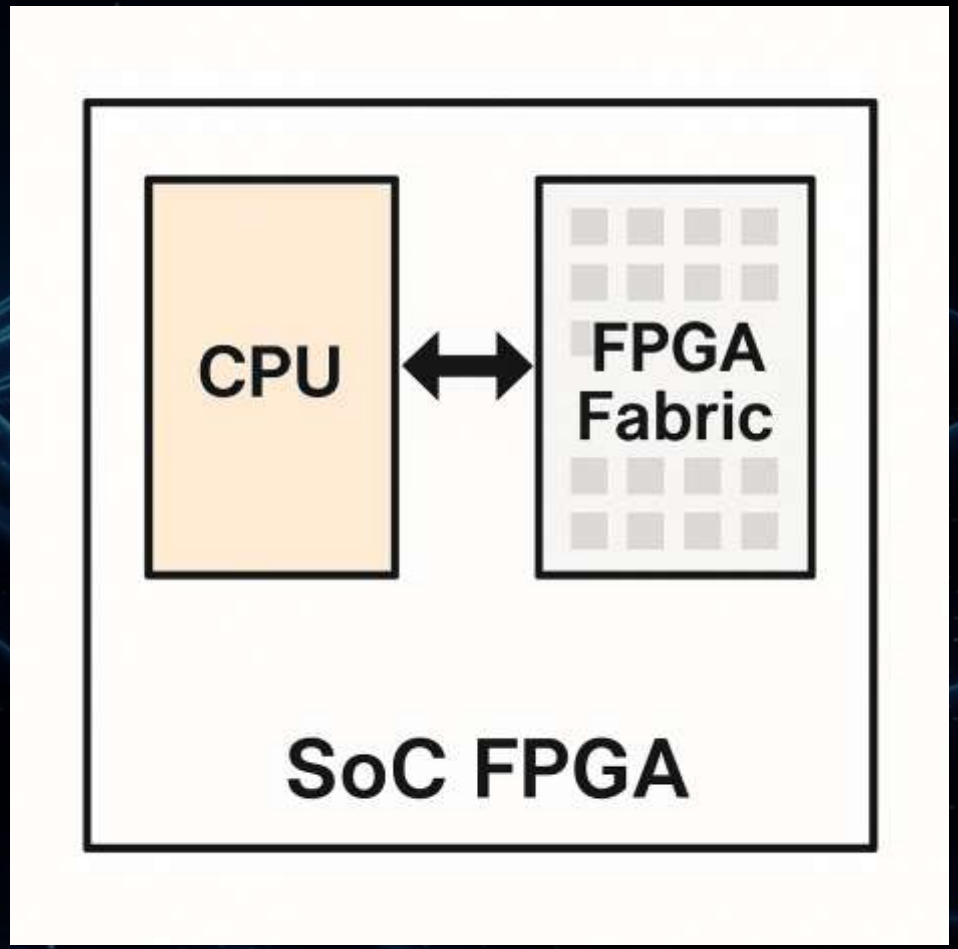

Real-World Attacks & Security Challenges

Real satellites have been hacked; ViaSat cyberattack exploited ground-segment vulnerabilities to gain privileges.

In-orbit reconfiguration and AI model updates introduce risks: performance degradation, service disruption and adversarial manipulation.

Lack of physical access amplifies risk.

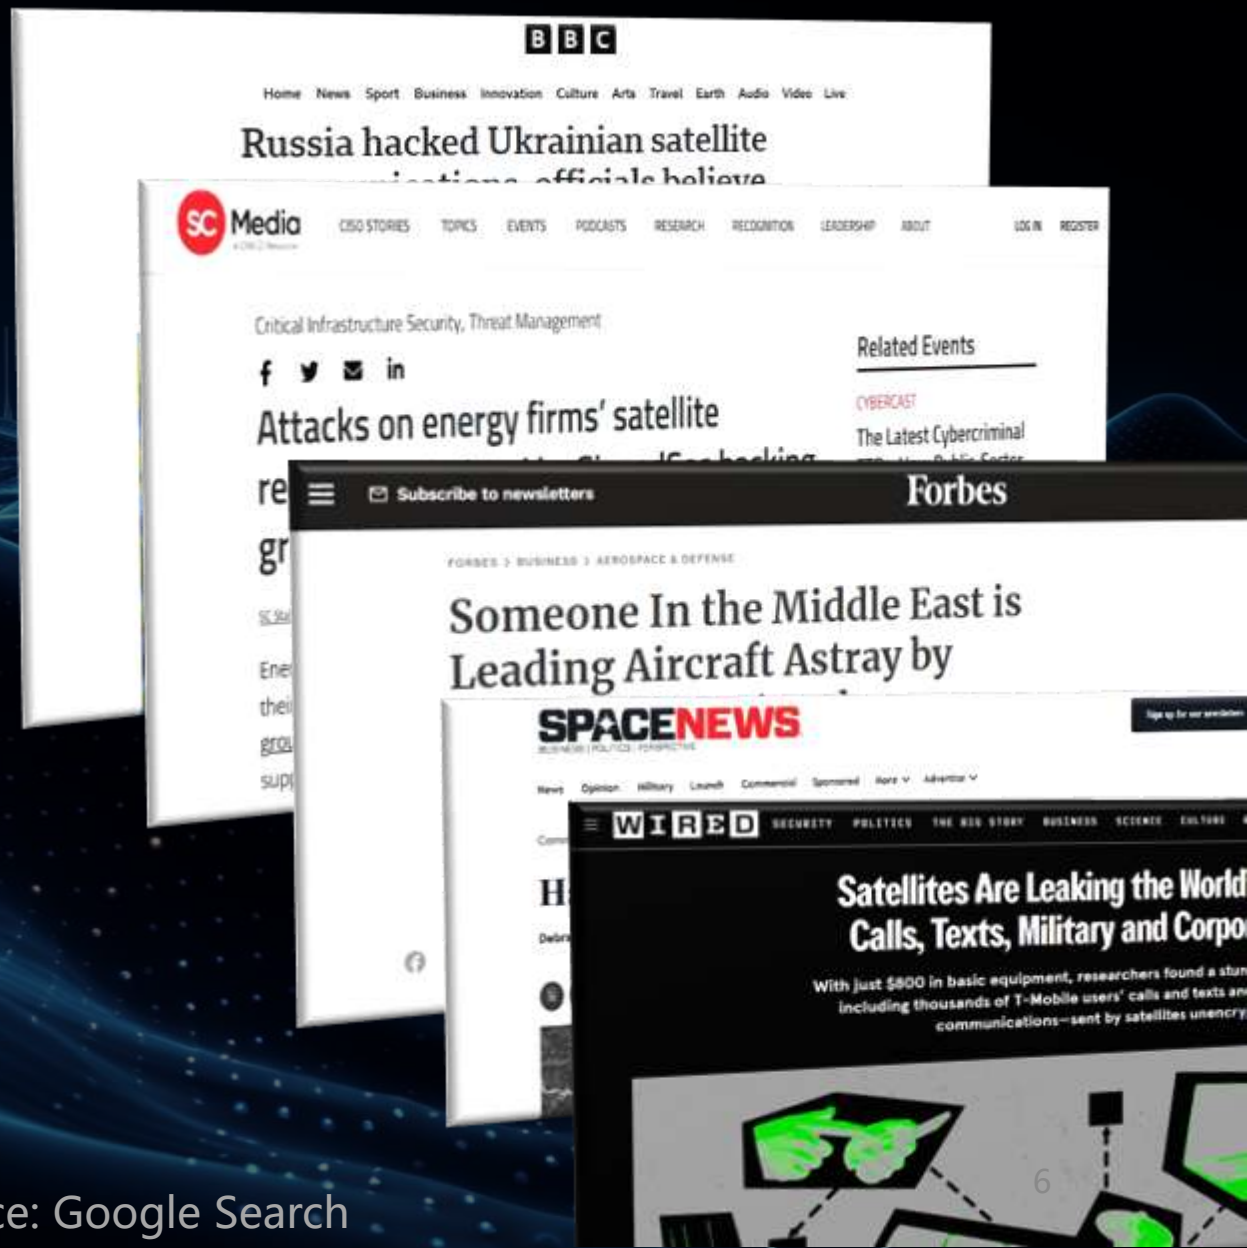

Source: Google Search

Fragmented Prior Work

1. Partial reconfiguration for cryptographic devices
2. lightweight reconfiguration security
3. onboard AI model updates
4. No unified end-to-end security framework for AI-enabled SoC FPGA satellites.

Existing Research Addresses Isolated Aspects

Need: integrated defense across boot, runtime and update phases.

Multi-Tenant Satellites & Federated Constellations

New multi-tenant model paradigms enable multiple stakeholders to share payloads and sensors.

Partitioning programmable logic into virtual FPGAs hasn’ t been adapted to satellites.

Federated constellations enable distributed learning, expanding the attack surface with AI model injection & exfiltration.

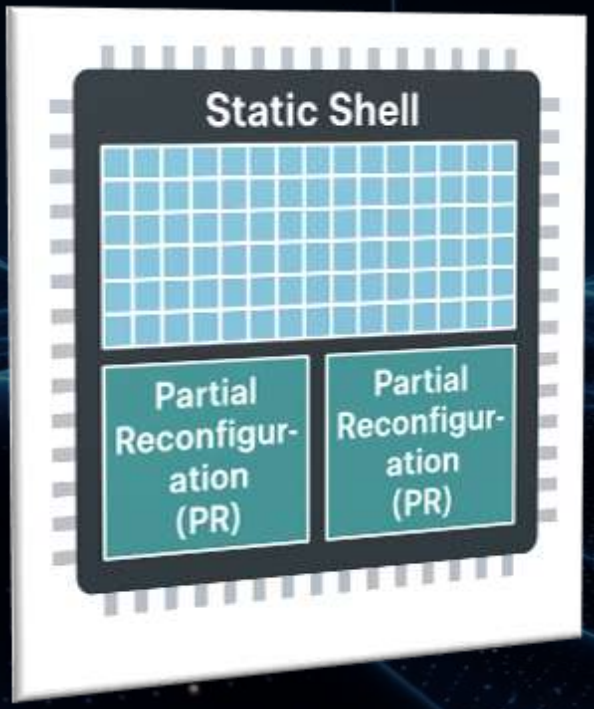

TABLE I
SURVEY OF SOC FPGA-BASED AI IMPLEMENTATIONS FOR AEROSPACE APPLICATIONS.

Authors	Year	AI Algorithms	SoC FPGA	Vendor	Applications
Pitsis et al. [28]	2019	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Space data classification
Ma et al. [29], [30]	2019	Autoencoder NN	Zynq UltraScale+ MPSoC	AMD Xilinx	Feature extraction and anomaly detection
Sabogal et al. [31]	2019	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Semantic segmentation of space imagery
Liu et al. [32]	2019	CNN	Arria 10 SX SoC	Intel Altera	Remote sensing image segmentation
Li et al. [33]	2019	SSD NN	Zynq 7000	AMD Xilinx	Remote sensing imagery analysis
Reiter et al. [34]	2020	BNN	Zynq 7000	AMD Xilinx	Real-time cloud detection
Lemaire et al. [35]	2020	BNN + CNN	Cyclone V	Intel Altera	Cloud classification and detection
Lent et al. [36]	2020	SNN	Zynq 7020	AMD Xilinx	Routing in space networks
Cosmas et al. [37]	2020	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Visual landmark recognition for navigation
Zhang et al. [38]	2021	YOLOv2 (CNN)	Zynq 7000	AMD Xilinx	Optical object detection
Rapuano et al. [1]	2021	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Cloud detection
Sabogal et al. [39]	2021	CNN	Zynq 7020, Zynq UltraScale+ MPSoC	AMD Xilinx	Semantic segmentation
Pacini et al. [40]	2021	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Real-time image classification
Pitonak et al. [41]	2022	CNN	Zynq 7020	AMD Xilinx	Cloud detection
Zhang et al. [42]	2022	GNN	Zynq UltraScale+ MPSoC	AMD Xilinx	SAR image classification
Abderrahmane et al. [43]	2022	SNN	Cyclone V	Intel Altera	Cloud detection
Papatheofanous et al. [44]	2022	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Satellite image segmentation
Perryman et al. [45]	2023	MobileNetV1, ResNet-50, GoogLeNet	XCVC1902 (VCK190)	AMD Xilinx	Edge computing in space
Ekblad et al. [46]	2023	YOLOv4-based NN	Zynq UltraScale+ MPSoC	AMD Xilinx	Autonomous navigation
Gao et al. [47]	2023	CNN	Zynq 7000	AMD Xilinx	CNN reliability evaluation
Carmeli et al. [48]	2023	SOM NN	Cyclone V	Intel Altera	Star pattern recognition
Coca et al. [49]	2023	ResNet	Zynq UltraScale+ MPSoC	AMD Xilinx	Burned area anomaly detection
Zhao et al. [50]	2023	YOLOv4-MobileNetv3	Zynq UltraScale+ MPSoC	AMD Xilinx	Object detection in satellite images
Mazouz et al. [51]	2024	YOLOv3	Zynq 7100	AMD Xilinx	Streaming object detection
Kim et al. [52]	2024	Reinforcement Learning	Zynq 7000	AMD Xilinx	Routing in LEO networks
Castelino et al. [53]	2024	Conv. Autoencoder	Zynq UltraScale+ MPSoC	AMD Xilinx	HSI artifact detection
Zhang et al. [54]	2024	Dehazing NN	Zynq 7000	AMD Xilinx	Image dehazing
Cratere et al. [55]	2024	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Cloud detection
Kim et al. [56]	2024	SqueezeNet	Zynq 7000	AMD Xilinx	Cloud detection
Li et al. [57]	2024	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Depth estimation in spacecraft
Upadhyay et al. [58]	2024	ResNetc	Zynq UltraScale+ MPSoC	AMD Xilinx	Cloud detection
Posso et al. [59]	2024	Mobile-URSONet	Zynq UltraScale+ MPSoC	AMD Xilinx	Pose estimation
Ciancarelli et al. [60]	2024	Autoencoders, CNNs	Xilinx ACAP	AMD Xilinx	Anomaly detection, SAR, RF
Leon et al. [61]	2024	UrsoNet, MobileNetV2, ResNet-50	Zynq UltraScale+ MPSoC	AMD Xilinx	Pose estimation and benchmarking
Barnwal et al. [62]	2024	CNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Galaxy classification
Bai et al. [63]	2024	CNN	Zynq 7020	AMD Xilinx	Particle identification
Jiang et al. [64]	2024	DNN	Zynq UltraScale+ MPSoC	AMD Xilinx	Hyperspectral anomaly detection
Shi et al. [65]	2024	CNN	Zynq-7000, UltraScale+ MPSoC	AMD Xilinx	Image classification
Justo et al. [66]	2024	CNN	Zynq 7030	AMD Xilinx	Hyperspectral segmentation
Renaut et al. [67]	2025	DNN	Zynq 7000	AMD Xilinx	Satellite pose estimation
Garcés-Socarrás et al. [68]	2025	CNN	VC190, Zynq UltraScale+ MPSoC	AMD Xilinx	Payload config and beamforming
Perryman et al. [69]	2025	CNN	XCVC1902 (VCK190), XCVE2802 (VEK280)	AMD Xilinx	Fault-tolerant AI acceleration

Threat Model

02

Satellite System & Attack Vectors

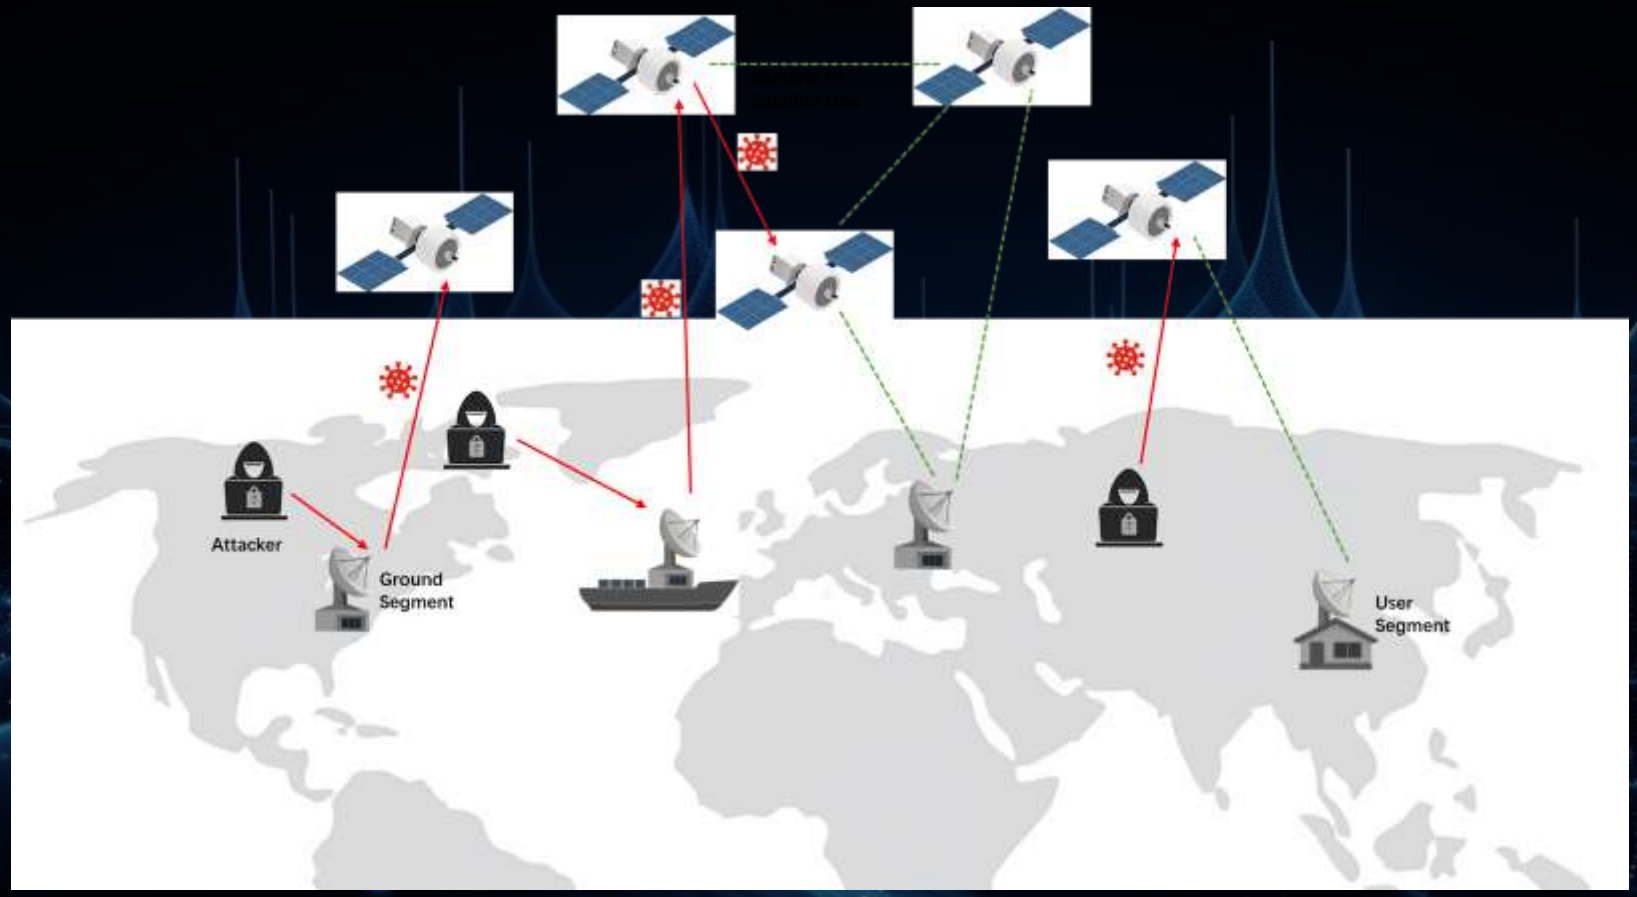

- A satellite system includes the space segment, ground segment and user segment.
- Adversaries can target ground infrastructure or satellites directly.
- A single compromised satellite can infect others via inter-satellite links.

Threat Model

01. Attackers

Have full knowledge of hardware & software, and legitimate SaaS access to upload custom bitstreams/models.

02. Other applications

Benign but not trusted; exploits can stem from software bugs, weak PS/PL isolation or poor authentication.

03. Goals

Include unauthorized data access, mission disruption and AI manipulation.

AegisSat Design

03

Defense-in-Depth & AegisSat Overview

A

Secure Boot and Root of Trust.

D

Robust rollback mechanisms for resilience.

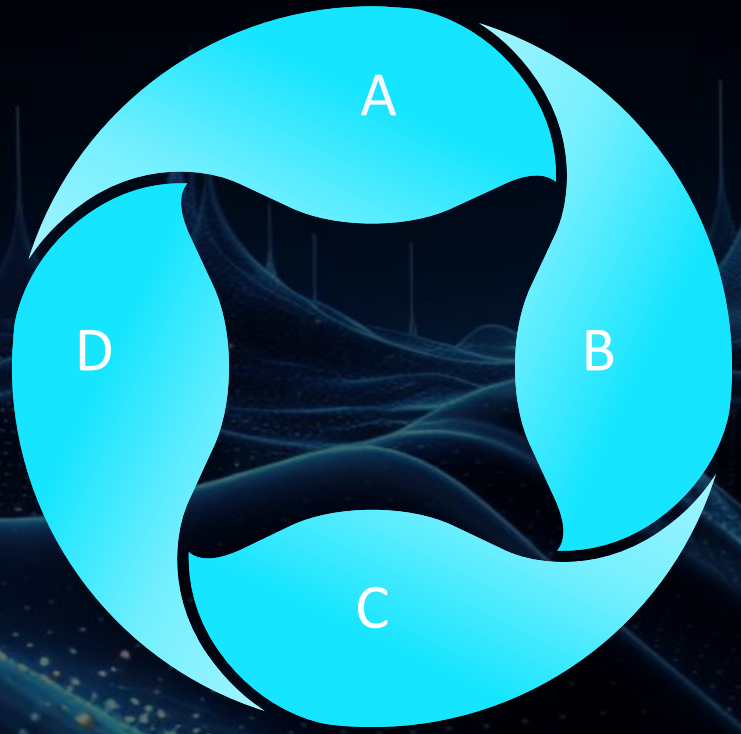

B

Trusted Execution and Hardware Isolation

C

Secure Update, Partial Reconfiguration, and AI Model Updates

Each layer reinforces the next, ensuring resilience even without physical access.

A Secure Boot and Root of Trust.

Chain-of-Trust

- Hierarchical chain of trust
- Cryptographic signatures and encryption protect against tampering.
- Only trusted firmware and AI models run.

Key Management

- Hashes of public keys
- Volatile keys
- Physical Unclonable Functions (PUFs)
- Tamper detection circuits

Bitstream Authentication

- Bitstreams are encrypted and signed
- Dedicated hardware handles decryption & authentication.
- Version control and anti-rollback mechanisms

Failure Handling

- Non-overwritable golden image, retry logic, safe-mode boot, and watchdog timers.
- Golden image stored in secure memory

B Trusted Execution and Hardware Isolation

Runtime Threat Landscape

- AI inference engines, OS services, and communication may contain flaws
- Malicious bitstreams
- Tight PS/PL coupling

Execution Isolation in PS

- TEE (Trusted Execution Environment)
- Firmware and AI Validation
- AI inference and other data handling
- MPUs and MMUs

Isolation of PL & Secure AI Co-Processing

- Hardware-level isolation
- Interface validation
- Secure AI co-processing
- Runtime monitors

Privilege Separation

- Layered isolation
- Localized resets
- Hardware-enforced limits
- Resilience mechanisms

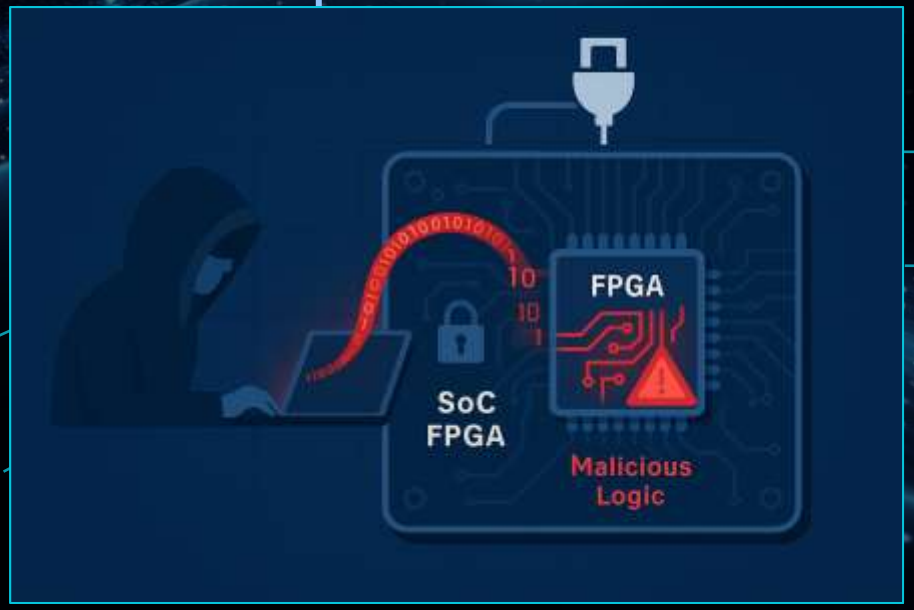

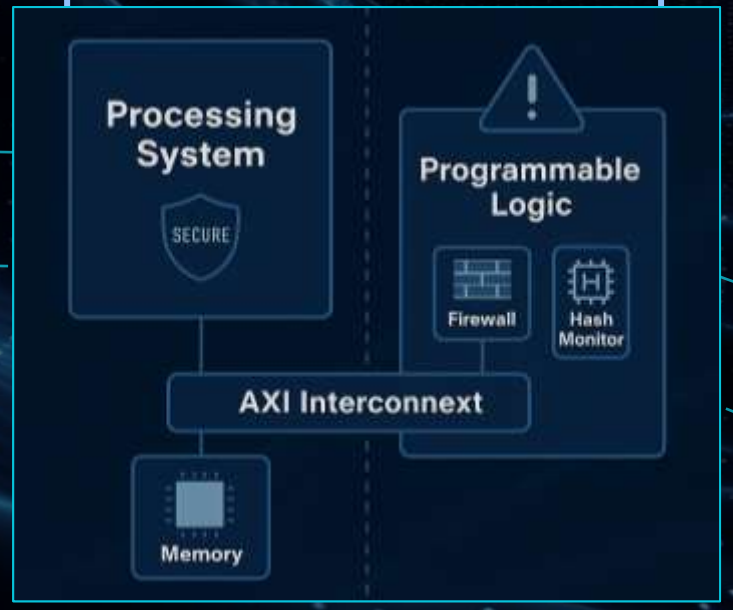

C Secure Update, Partial Reconfiguration, and AI Model Updates.

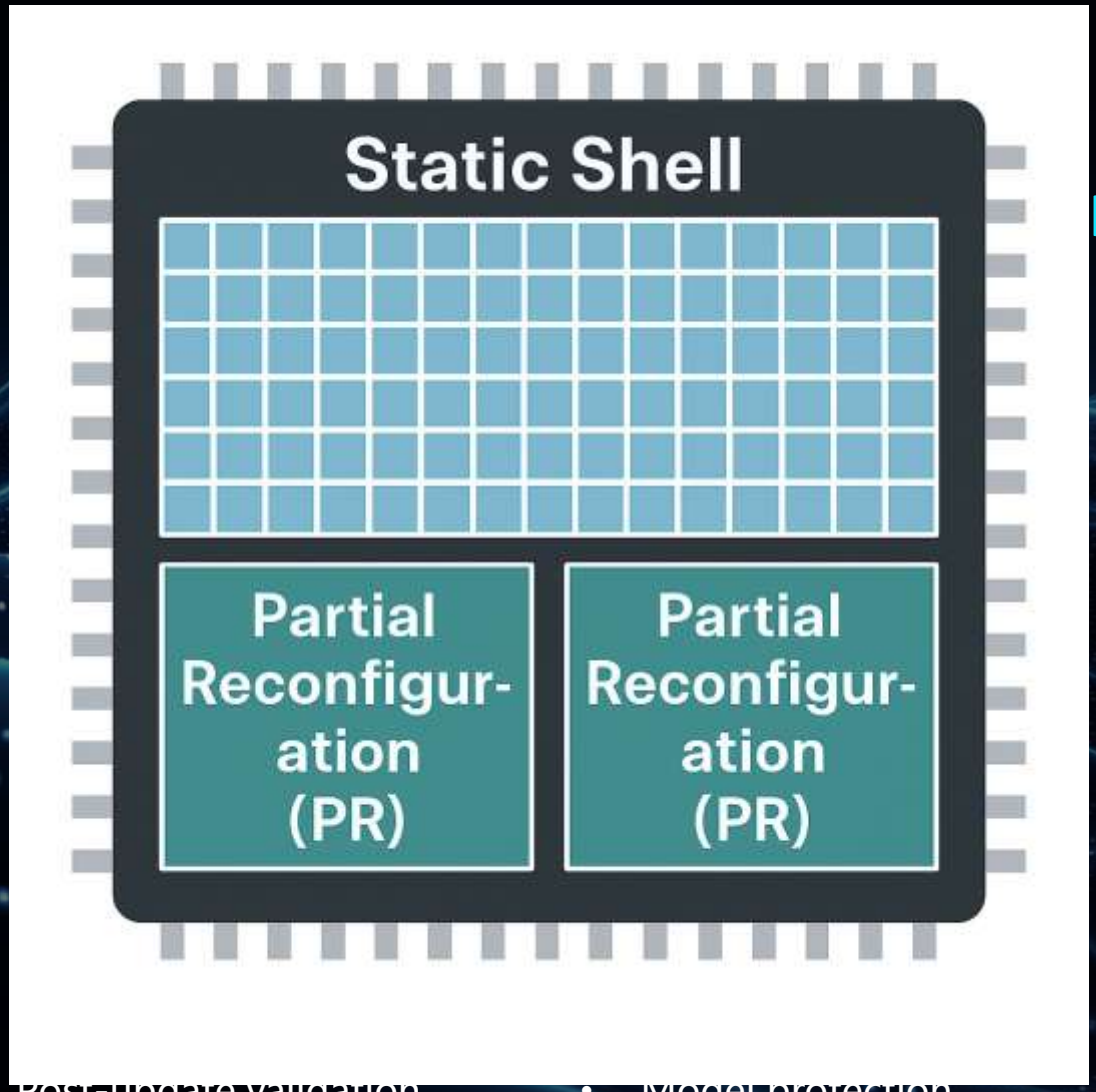

- Authenticated transmission
- Staged validation

- Post-update validation
- Active threat detection

- Model protection
- Federated security

Operational &

- Safe scheduling
- Reliable execution

D Robust rollback Mechanisms

AegisSat provides robust rollback mechanisms that restore a known-good state from a golden image, ensuring service continuity and preventing persistent compromise.

Implementation

04

Implementation Platform: ZCU102 Setup

Proof-of-concept implemented on Xilinx ZCU102 board: quad-core ARM Cortex-A53 with TrustZone, dual Cortex-R5 and FPGA fabric.

Source: https://www.amd.com

Partitioning into vFPGAs & Trust Anchor

Programmable logic partitioned into one static region (FPGA shell) and two dynamic vFPGA partitions. Each vFPGA has dedicated I/O and memory ranges.

In the processing system, one app, named FPGA Trust Anchor (TAFPGA), runs in the Secure World; user apps run in the Normal World.
Inter-core communication via Xilinx Inter-Processor Interrupts and shared memory.

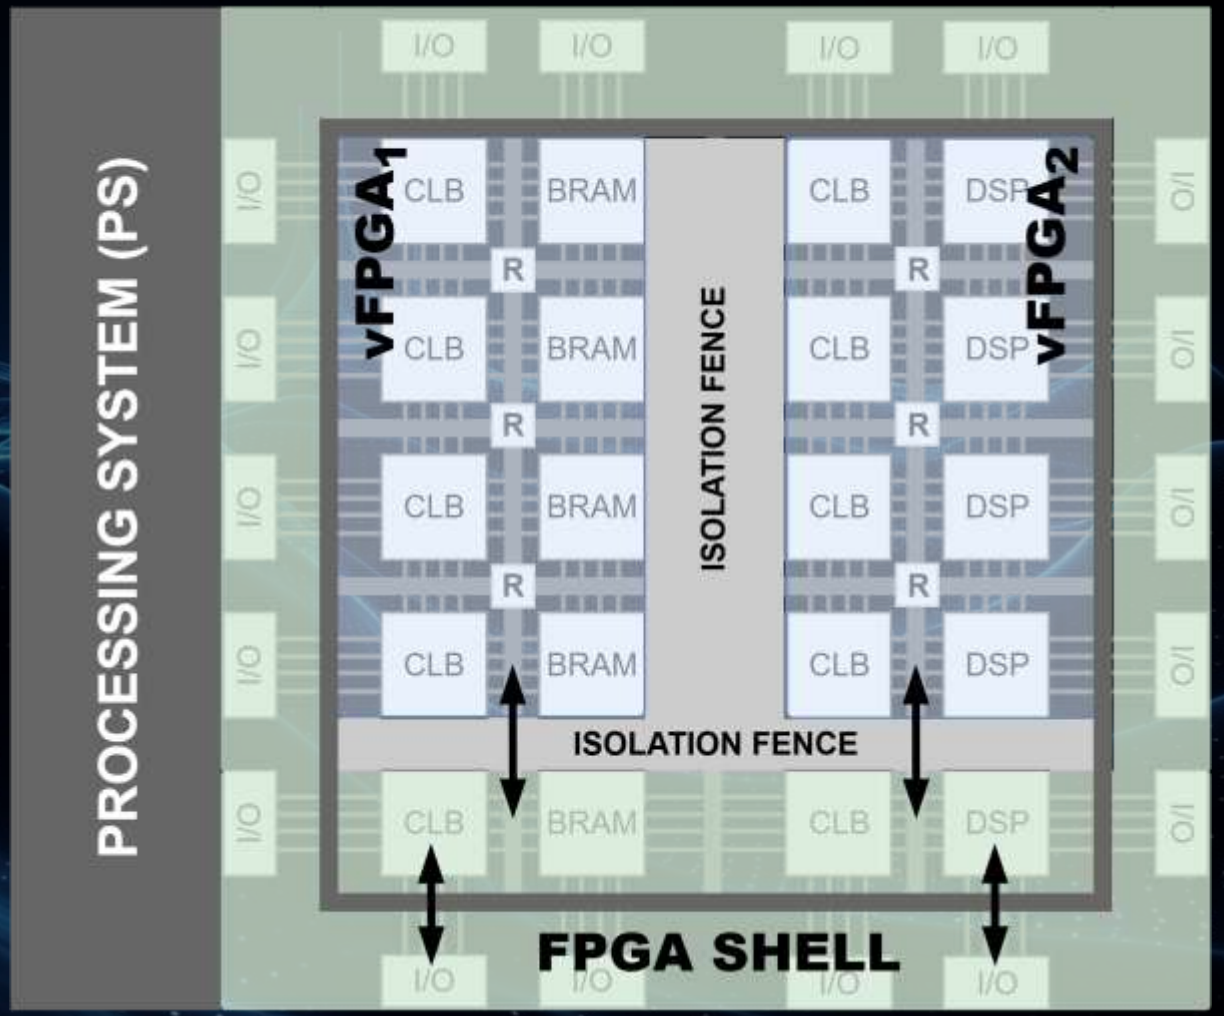

Secure Reconfiguration Workflow (Steps)

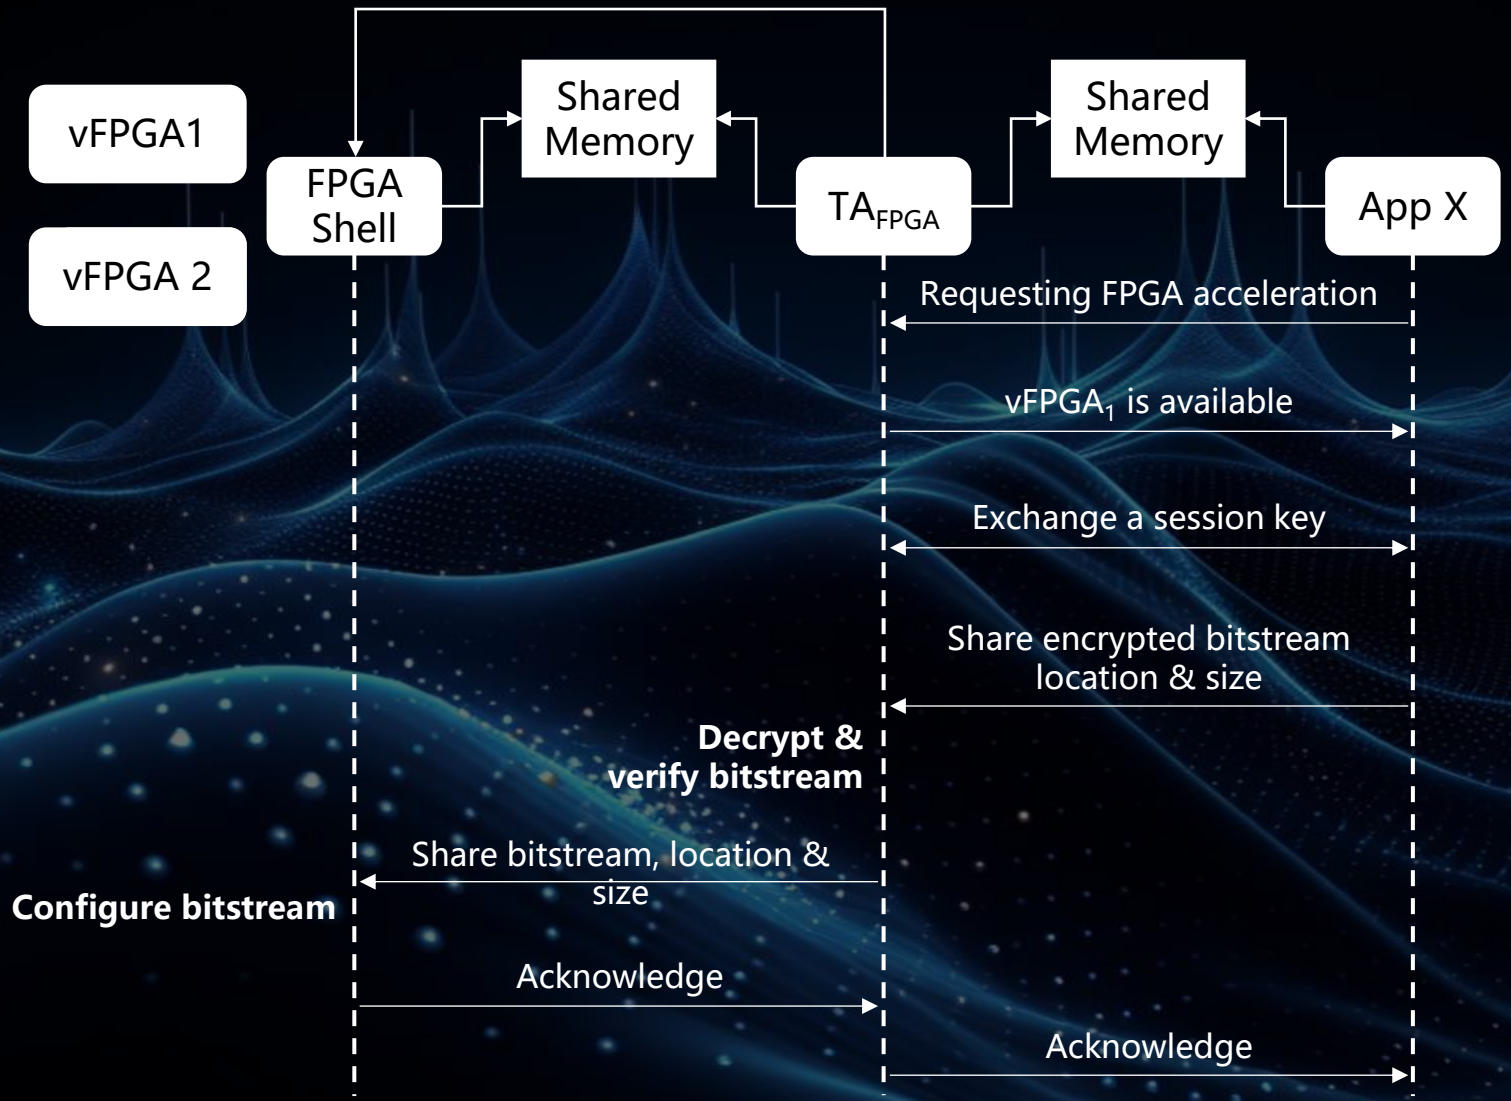

Implementation Results: Resource & Performance

vFPGA1: Lightweight CNN accelerator: 3×3 convolution, quantization, ReLU, pooling on 6×6 input.

vFPGA2: Shift circuit.

Mean partial reconfiguration: vFPGA1 ~495.21 ms (σ≈8.64 ms), vFPGA2 ~528.21 ms (σ≈0.27 ms).

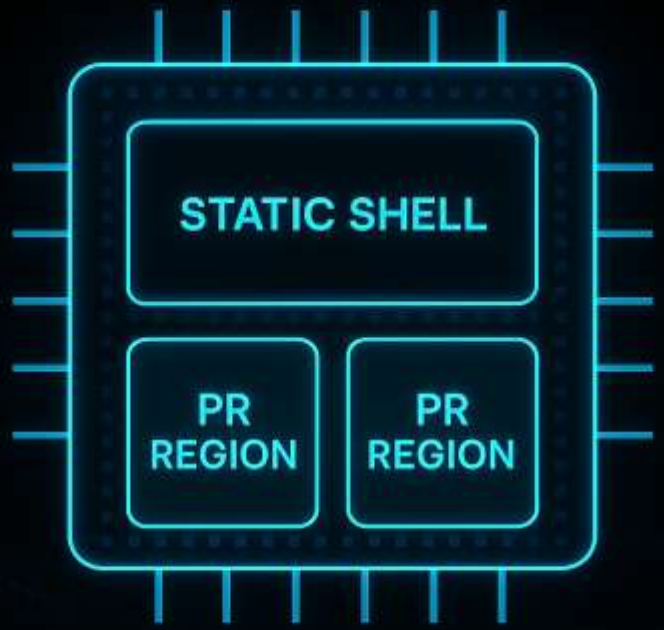

Research Outlook

05

Research Outlook: Post-Quantum & Autonomous Attestation

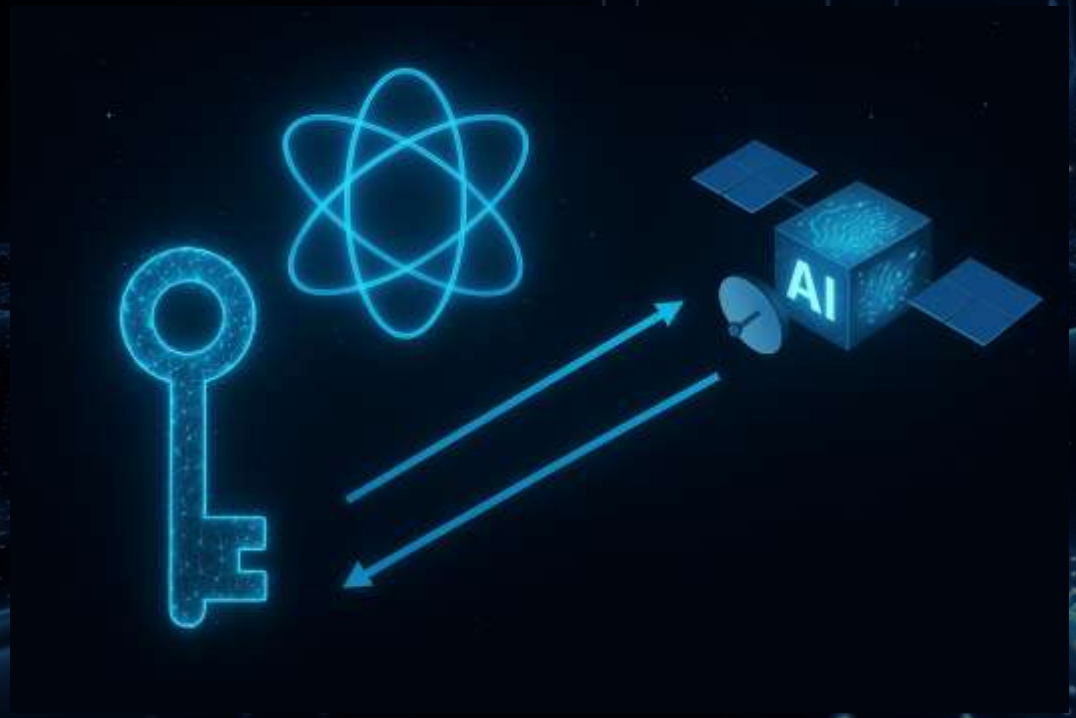

A Post-Quantum Cryptography

Classical public-key cryptography, such as RSA/ECDSA are vulnerable to quantum attacks; need to integrate PQC schemes (e.g., CRYSTALS-Dilithium, SPHINCS+) into SoC FPGA toolchains.

B Autonomous Attestation

Long-duration deep-space missions require continuous trust evaluation without ground oversight.
Need lightweight, self-verifying architectures for in-orbit measurement of PL fabric under communication delays.

Research Outlook: Energy-Constrained Isolation & Cross-Domain Co-Design

C Energy-Constrained Isolation

TrustZone, AXI firewalls and SMMUs add latency and power overhead; need optimized RTL-level techniques (logic-privilege tags, dynamic bus gating) and energy-adaptive security policies.

D Cross-Domain HW/SW Co-Design

Security must span AI model lifecycle – synthesis, compiler output validation, key provisioning and telemetry. Open-source toolchains could balance security and accessibility. Integrating hardware-assisted side-channel monitors (glitch sensors, power anomaly detectors) within FPGA fabric is underexplored.

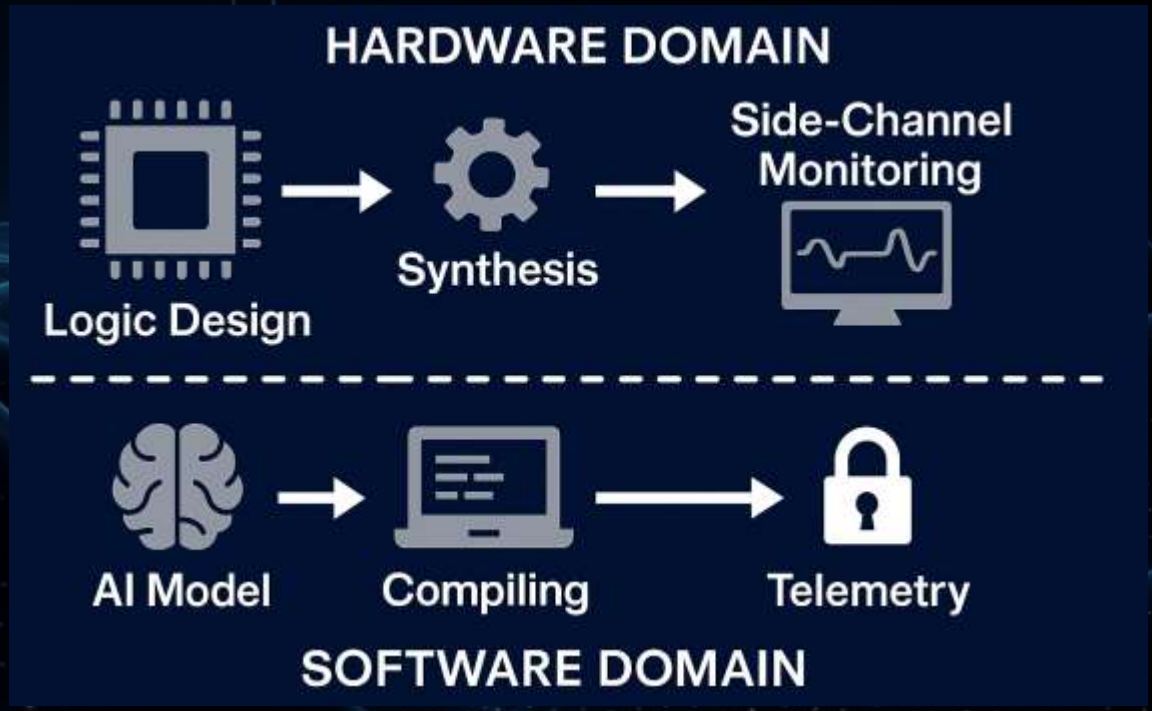

Research Outlook: Federated Learning

E

Federated Learning in Space

Federated Learning in Space: Constellations of cooperating satellites perform federated AI; secure aggregation protocols exist but hardware-rooted defense against poisoning/Byzantine faults across hundreds of nodes remains open.

Conclusions

06

Conclusion

AegisSat provides a layered security framework for AI-enabled SoC FPGA satellites, combining secure boot, isolation, authenticated updates and rollback protection.

Proof-of-concept demonstrates feasibility on modern SoC FPGA devices with acceptable overhead.

Future work: integrate PQC, develop autonomous attestation, optimize energy-aware isolation, advance cross-domain co-design and secure federated learning.

Thank You!