Microsoft

Secure Management of Hyperscale Cloud Network Accelerators

Presenter: Faye Yang Xiling Sun

11/7/2024 WISH'24

Microsoft

Cloud Attacks

*Which will cause **data breaches, data loss, unauthorized access, disruption to the services** and significant **operational challenges**...*

Some famous DDoS attacks:

- Oct 2023: Google claimed the largest DDoS attack that peaked at 398 million RPS.
- Nov 2021: Azure experienced the DDoS attack reached a throughput of 3.47Tbps.
- Feb 2020: AWS reported a massive DDoS attack at a rate of 2.3 Tbps.

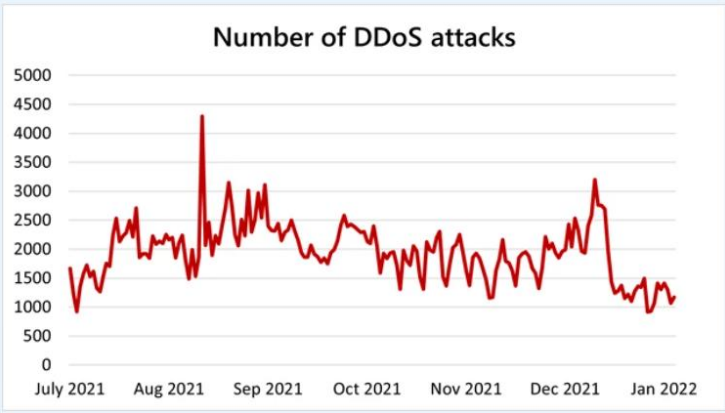

Source: https://azure.microsoft.com/en-us/blog/azure-ddos-protection-2021-q3-and-q4-ddos-attack-trends/

Cloud attack types

Source: Top 10 Cloud Attacks and What You Can Do About Them - Aqua (aquasec.com)

Microsoft

Why secure management of Network Accelerator Cards (NAC) matters?

- Secure platform (secure by design) is the foundation of modern data center infrastructure.
- Secured management of NAC ensures resilient against threats and vulnerabilities.
- Strike a balance between security and usability for cloud networks.

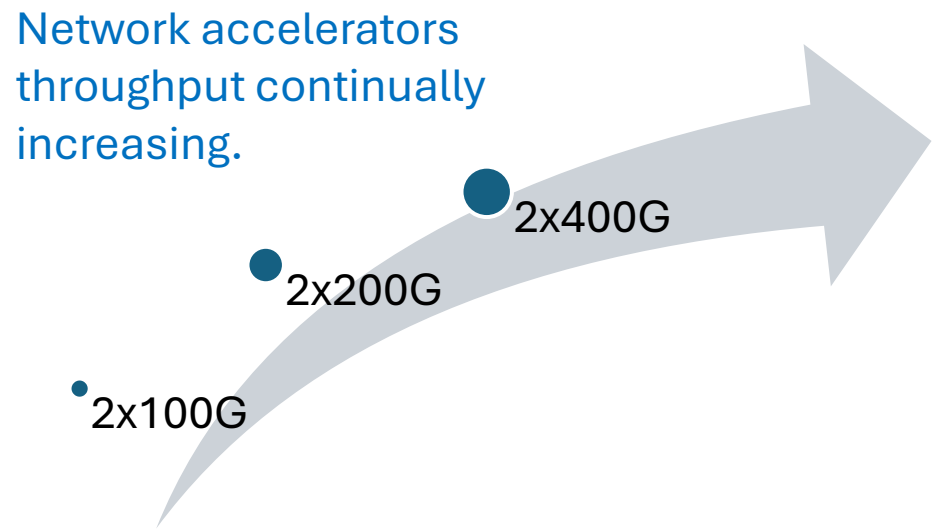

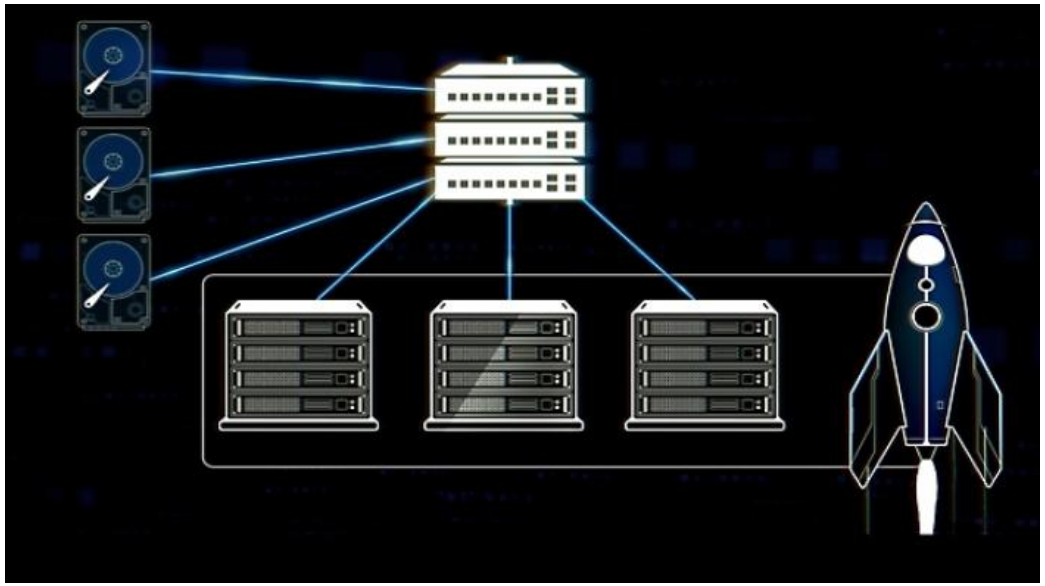

Azure boost: https://azure.microsoft.com/en-us/products/virtual-machines/boost/

Microsoft

Design Principles of Next-Gen Secure NAC Management

Advanced Hardware Interfaces

- Changing USB to SGMII
- Adding I3C interface

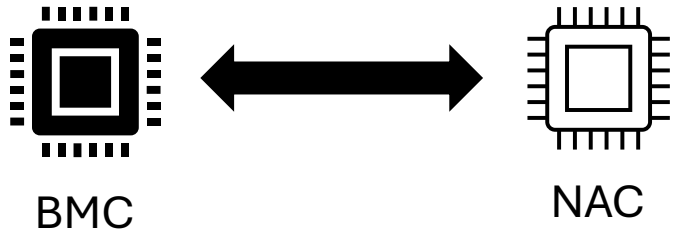

Layered Security Strategies

- Enforcing security from device SoC level
- Strengthening platform integrity via attestation
- Hardening remote access protocols with certificated-based authentication

- Remote management
- Hardware monitoring
- Firmware update and recovery

Microsoft

Main interfaces of NAC

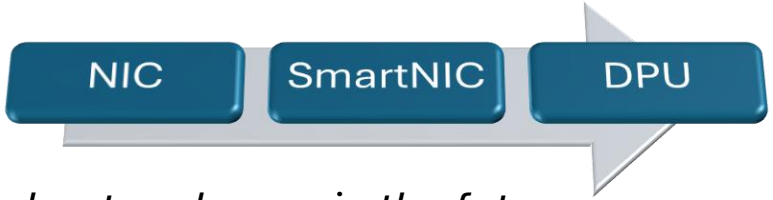

Which offloads networking, storage, management tasks from the host and more in the future…

Interfaces	Functions
UART	Serial port providing out of band admin console to the SoC on the NAC; debugging and diagnostics
USB 2.0 (480Mbps)	Storage allowing BMC/RM to transfer files to and from the NAC; USBvNIC communication for diag/telemetry/ config files transfer etc.
I2C	RoT communication; telemetry; OOB update and reset
Ethernet	Network connectivity; data transfer; traffic management
PCIe	Data plane path and control
1Gb (BMC to RM)	Out of band control; telemetry

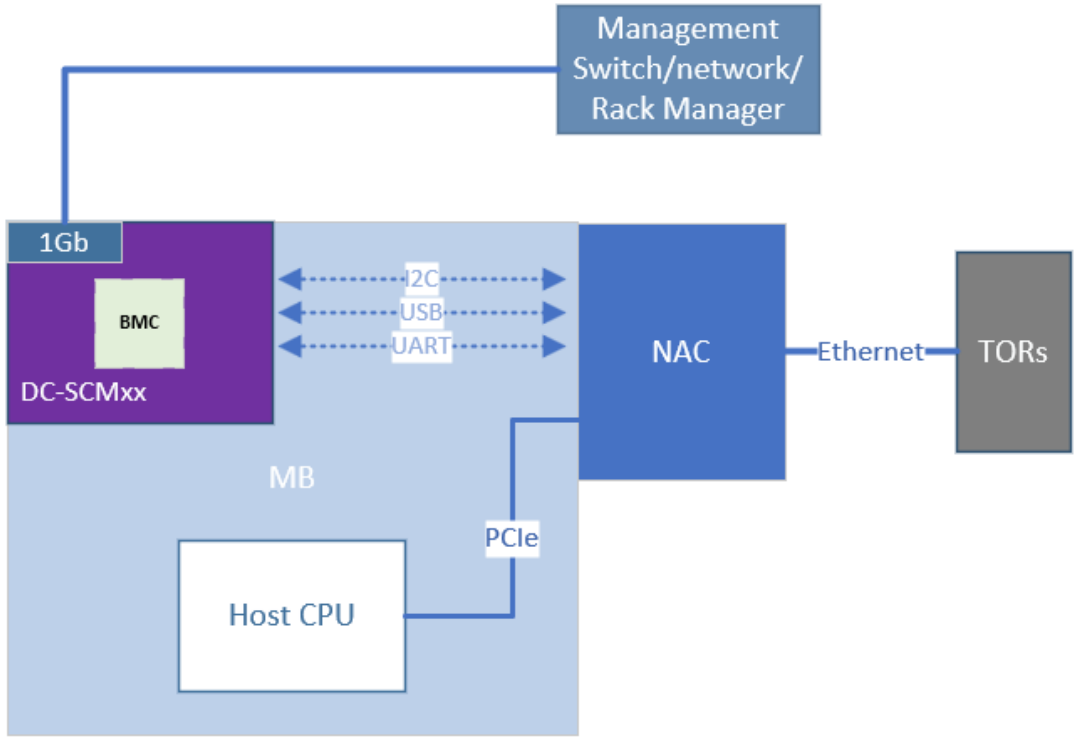

System view of simplified connectivity of Network accelerators

Microsoft

Changing USB 2.0 to SGMII

- Improved speed and SI (signal integrity)
  - 1.0 Gbps vs. 480Mbps
  - Better noise immunity and SI
- Protocol overhead and complexity
  - SGMII is simpler and more direct
  - Reduce the potential vulnerabilities
- Data transfer security
  - Reducing data corruption and interception
- More flexibility to system design
  - Not all servers support USB
  - Better performance as high bandwidth and low latency
  - Better support **1 to N** (NAC to hosts) or **N to 1** design (NACs to host)

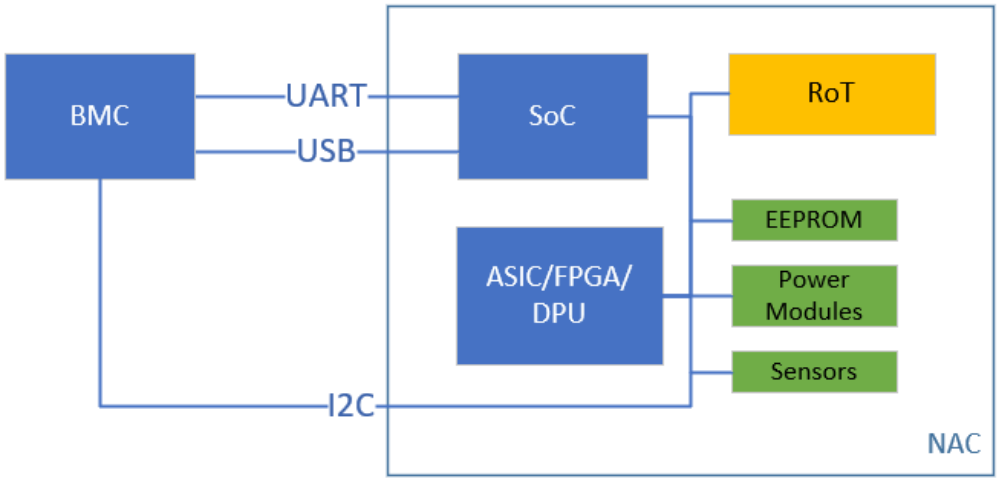

NAC-BMC interfaces on management platform

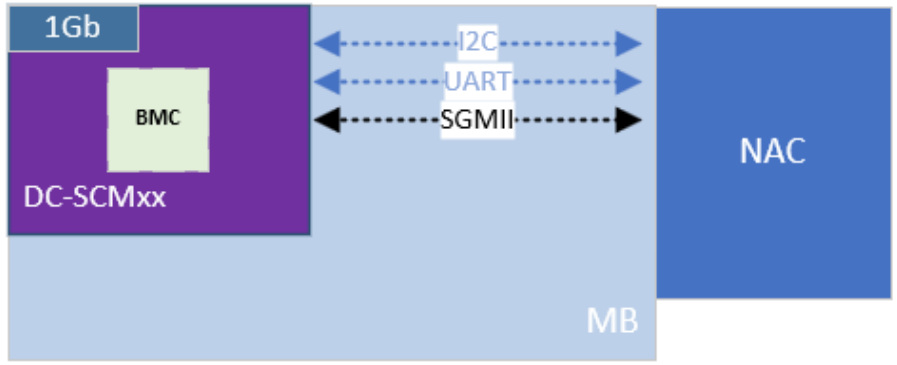

Microsoft

Adding I3C interface

- Speed improvement
  - The speed of I2C begins to show limitations for these OOB controls
  - I3C can achieve up to 12.5MHz in SDR mode
- Dynamic addressing
  - More flexibility for NAC integration into different server platforms
- Hot-join
  - Less interruption and more reliability
- In-band Interrupt
- Power Efficiency

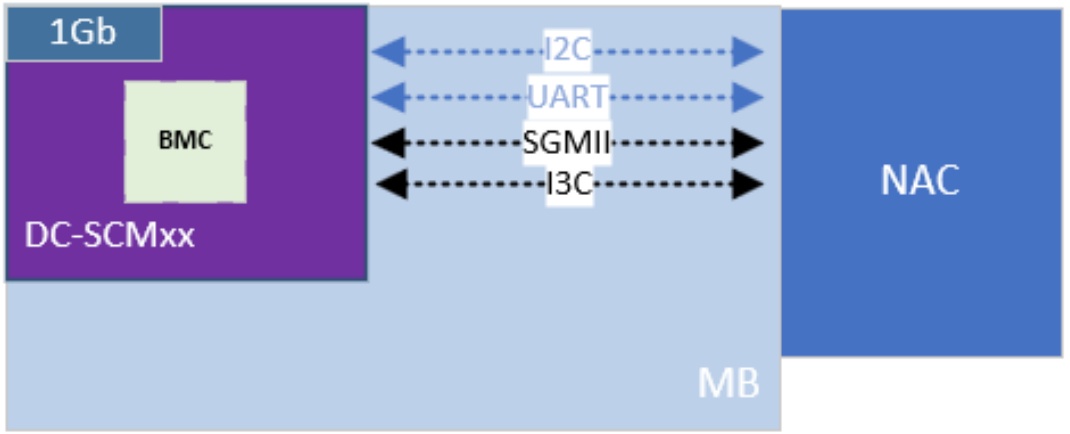

Microsoft

Layered Security Strategies

- Enforcing security from device SoC level
  - Root-of-Trust (RoT) subsystem integration
  - Trusted Execution Environment (TEE) establishment
- Strengthening platform integrity via attestation
- Hardening remote access protocols with certificated-based authentication

Microsoft

RoT Security Subsystem Overview

- A foundational security component embedded in a device's hardware.
- Provides a secure and immutable basis for all cryptographic operations.
- Compliant with NIST 800-193 platform firmware resiliency guidelines
  - Protection: Secure update
  - Detection: Secure Boot, Measured boot, Attestation
  - Recovery: Obtain operational images from a trustworthy entity

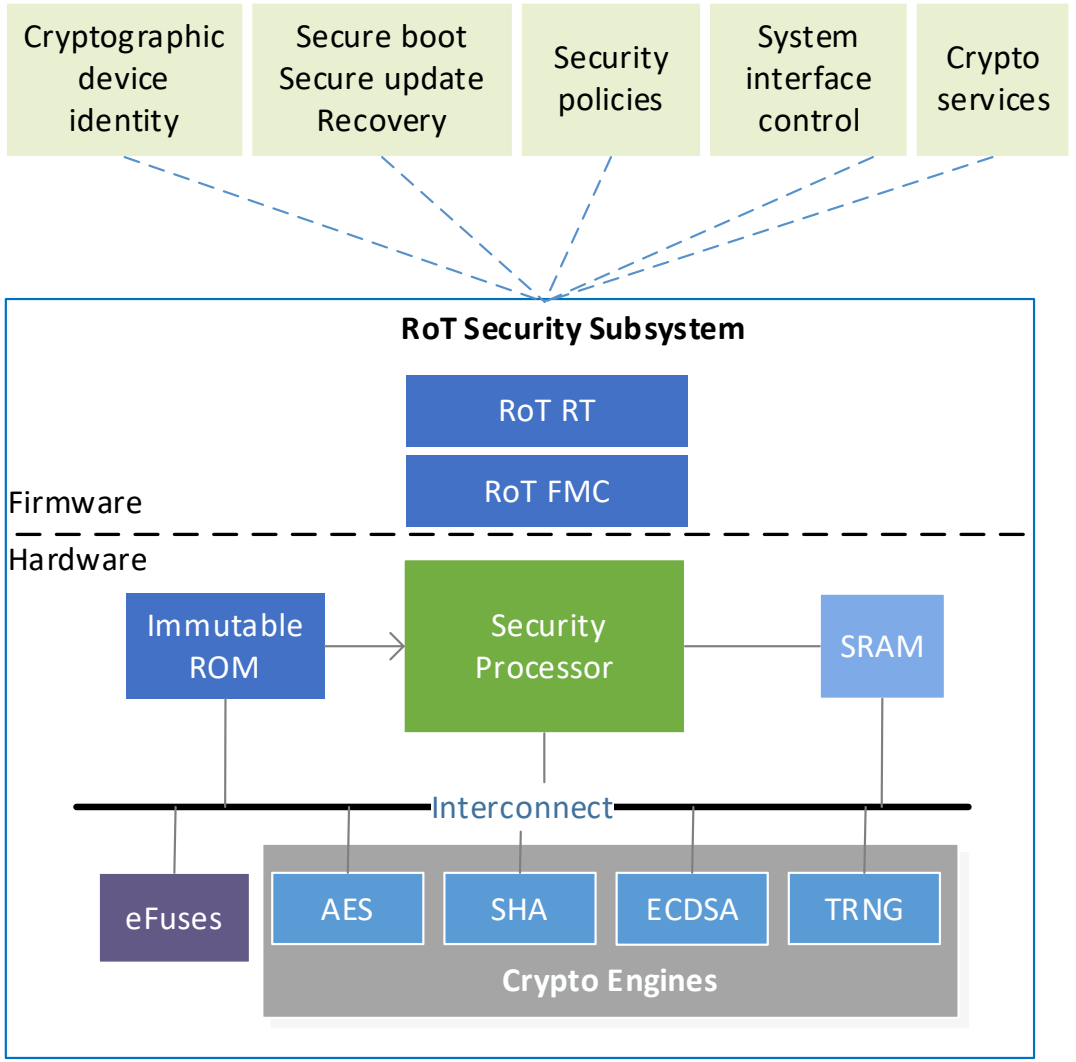

RoT security subsystem architectural view

Microsoft

Hardware-based Device Identity

Unique Device Secret (UDS)

- A cryptographic key stored in the device's hardware.
- Unique to each device, used to generate a device identity.

Device Identifier Composition Engine (DICE)

- A standard developed by the Trusted Computing Group (TCG).
- Establishes a chain of trust from the hardware to the software, ensuring that each layer is verified and secure.
- Creates unique, verifiable identity certificates for the device and its booted firmware.
- Small footprints, suitable to implement via HRoT in embedded devices.

Microsoft

DICE in a Nutshell

DeviceID Key

Derives an ECC key pair that will be same for as long as the Layer 0 firmware stays the same.

Alias Key

Derives a second key pair, depends on the identity of the Layer 1 and is same for as long as L1 stays the same.

Alias Key Certificate

Creates a certificate for the Alias Key using the Device ID private key.

Certificate Signing Request (CSR)

Creates a certificate signing request to simplify manufacturing flows where the device is vendor certified.

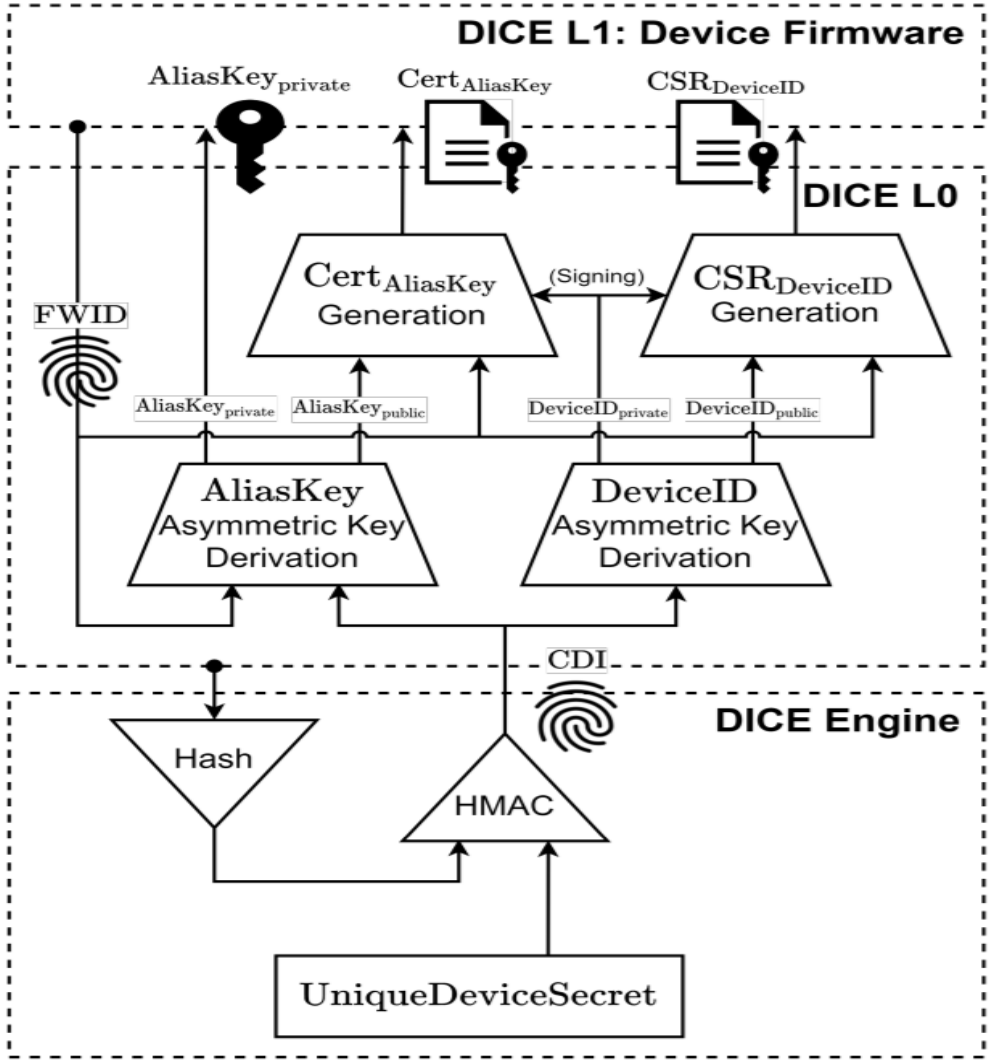

source: https://www.microsoft.com/en-us/research/uploads/prod/2021/02/main.pdf

Microsoft

Security-Centered NAC SoC Reference Architecture

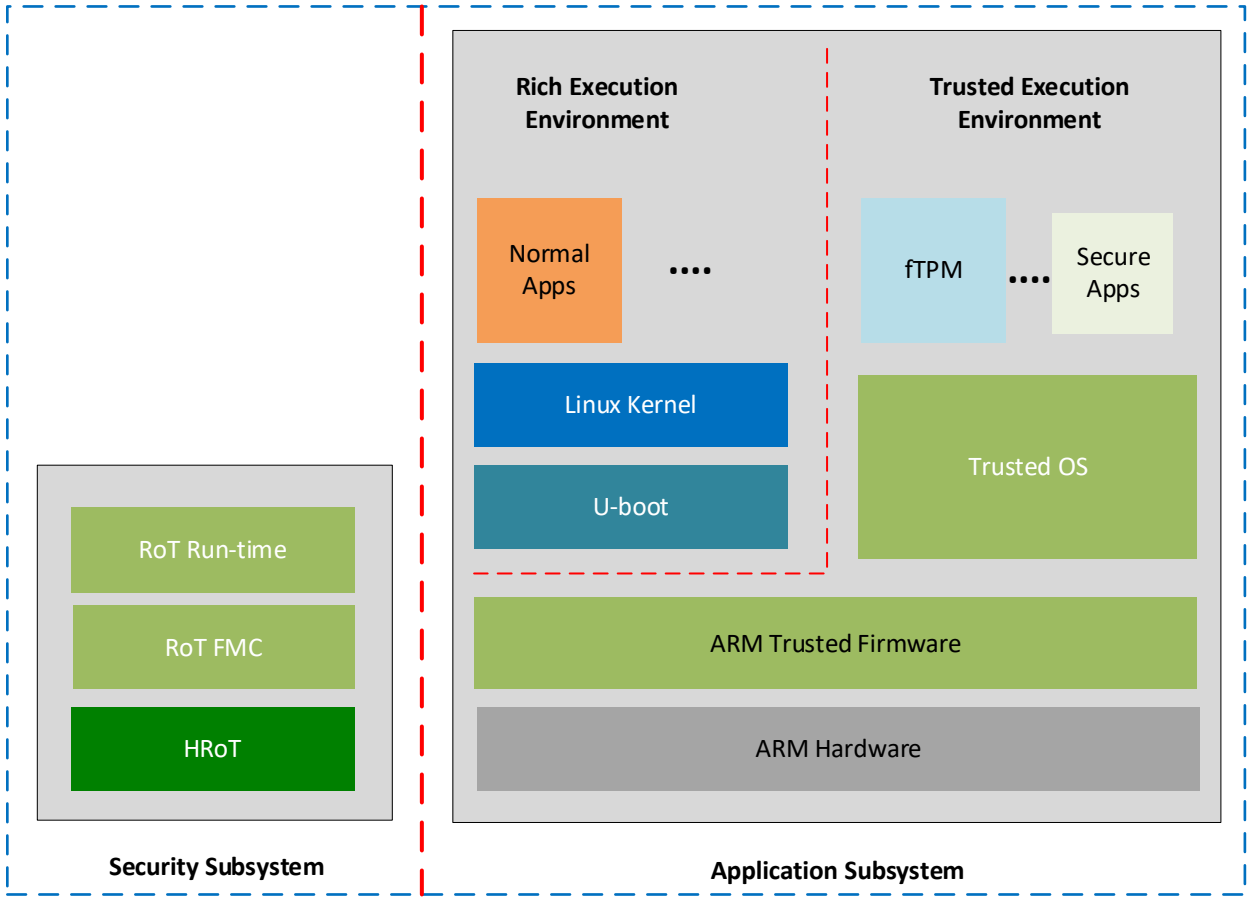

- Secure by Design
  - Hardware isolation
  - Integrated hardware RoT
  - RoT firmware stack compliant with NIST 800-193 guidelines.
  - Adopt Arm® TrustZone® technology to establish trusted execution environment(TEE).
  - Integrate firmware-based TPM (fTPM) as trusted application (TA).
- Approaches of RoT subsystem integration
  - Vendor-specific RoT implementation
  - Open-source Caliptra RoT integration

Microsoft

Strengthening Platform Integrity: Attestation via Platform RoT

Concept of Attestation

- Process of dynamically establishing and verifying the trust in device.
- Verifier establishes trust in the device by
  - Authenticating Device's Alias certificate chain
  - Verifying the signed firmware measurements with a challenge-response.
- Standardized attestation protocols
  - Security Protocols and Data Models (SPDM)
  - Project Cerberus Challenge Protocol

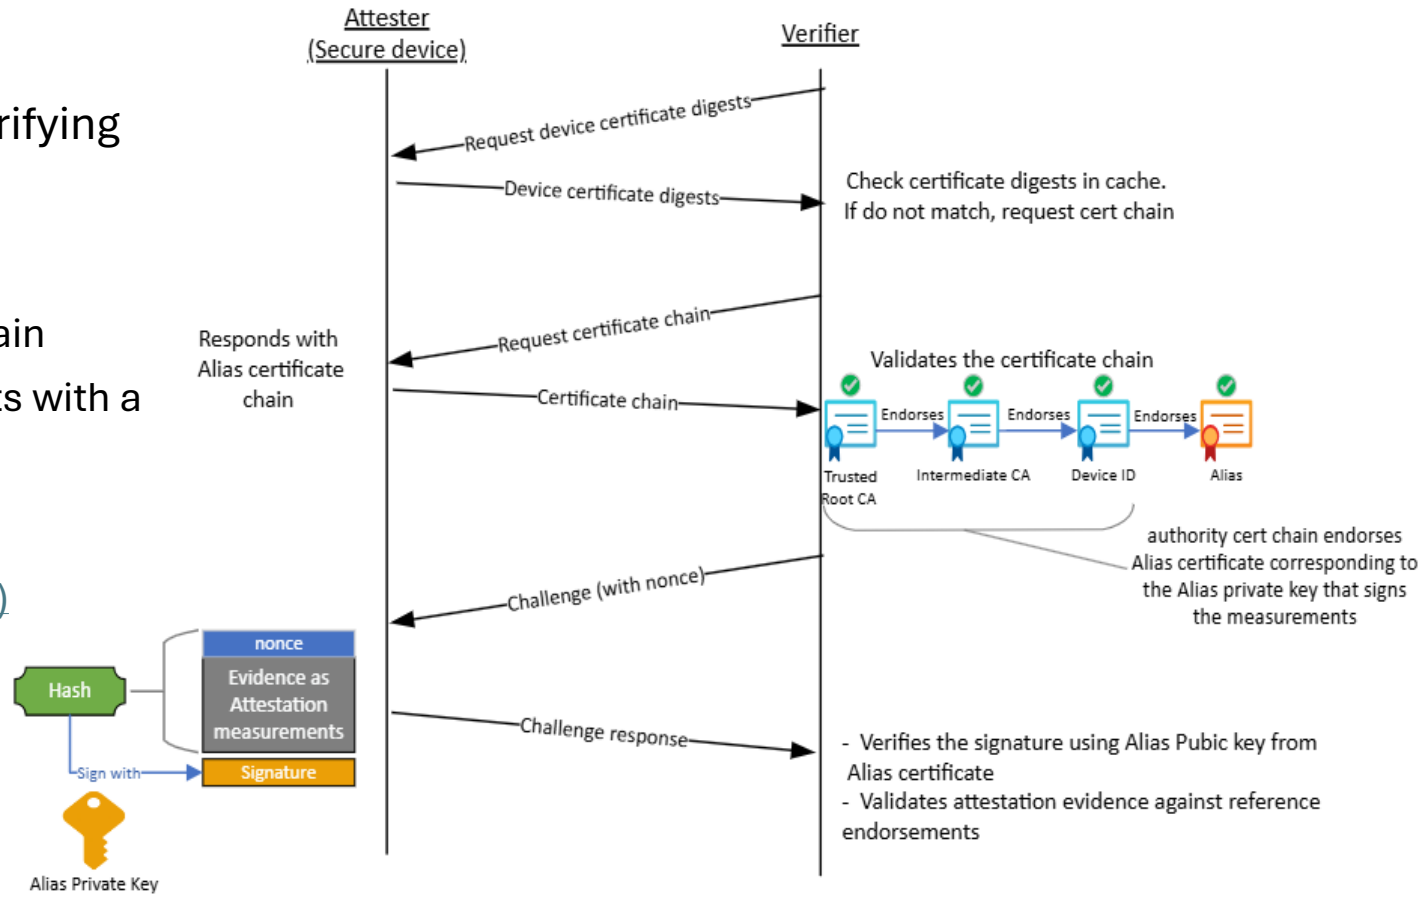

Microsoft

Platform Attestation Flow

1. **Validate the identity**
   - Platform RoT validates the device identity of the AC-RoT using the device's identity certificate chain.
2. **Authentication challenge**
   - Confirms that each device can use the private authentication key linked to the verified certificate chain by issuing a challenge to each device.
3. **Measurement verification**
   - Platform RoT extracts the digitally signed measurements reported by the AC-RoT and verifies them against reference measurements.
4. **Acceptance or remediation**
   - Based on the verification results, the Platform RoT either accepts the device or initiates a remedial action.

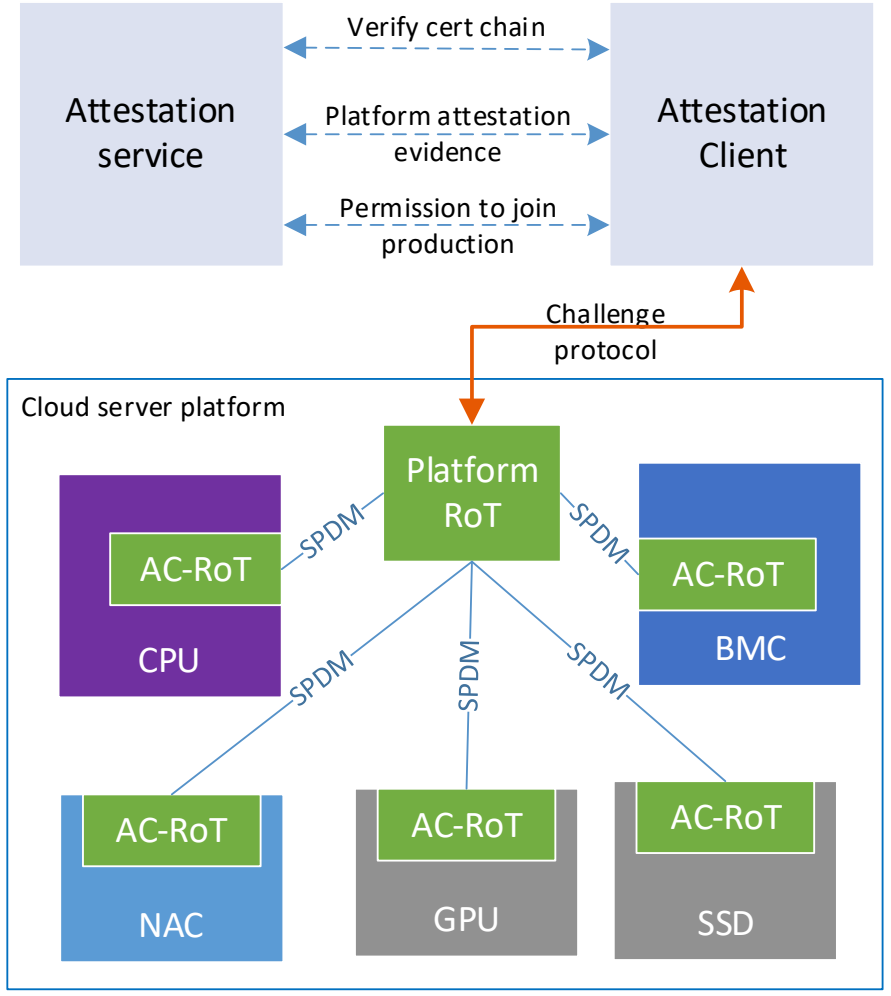

Simplified platform attestation workflow

Microsoft

Hardening Remote Access Protocols with Robust Certificate Authentication

Limitation of SSH public-key based authentication

- **Scalability issues**: Challenges in efficiently managing keys and connections.
- **Key management challenges**: Complexity in distributing, rotating, and revoking keys. Difficulties in tracking and securing a large number of keys, leading to potential vulnerabilities.
- **Insecure Trust-On-First-Use (TOFU) model**: trust in Server is established the first time a connection is made, which may pose risks if the initial key is compromised.

Benefit of SSH certificate-based authentication

- **Scalability**: A trusted CA server can be scaled up to securely issue certificates to designated users and devices.
- **Streamlined key management**: Eliminates the need for key distribution, centralizes security policy application.
- **Enhanced security**
  - **Access control**: Manage the accessibility and lifetime of certificates to mitigate unauthorized access risks.
  - **Granular permissions**: Configure certificate attributes to define access permissions with precision, ensuring users have only the necessary access.
- **Automated mutual authentication**: Host and client authenticate each other automatically, eliminating the need for passwords.

Additional security measures are still needed to protect the private keys

Microsoft

Prototype: OpenSSH RoT-backed Certificate Authentication on BMC

- SSH host certificate
  - Certificate based authentication for host.
  - Certificate restrictions can limit access by validity period, user, command, IP address, and more.
  - Leaf SSH cert rooted to hardware root CA
- BMC RoT subsystem
  - Generate a random SSH Host key pair on every boot or use a persistent key pair.
  - SSH host private key remans in RoT subsystem
  - Isolated from Linux and ARM64.
  - Handle the signing requests with the requested key slot.
- SSH host authentication
  - Native SSH service redirects signing requests to IRoT via PKCS11 provider.
  - Measurement and attestation info available in certificate chain
  - Solves TOFU issue.
  - Key revocation/renewal based on certificate policy.

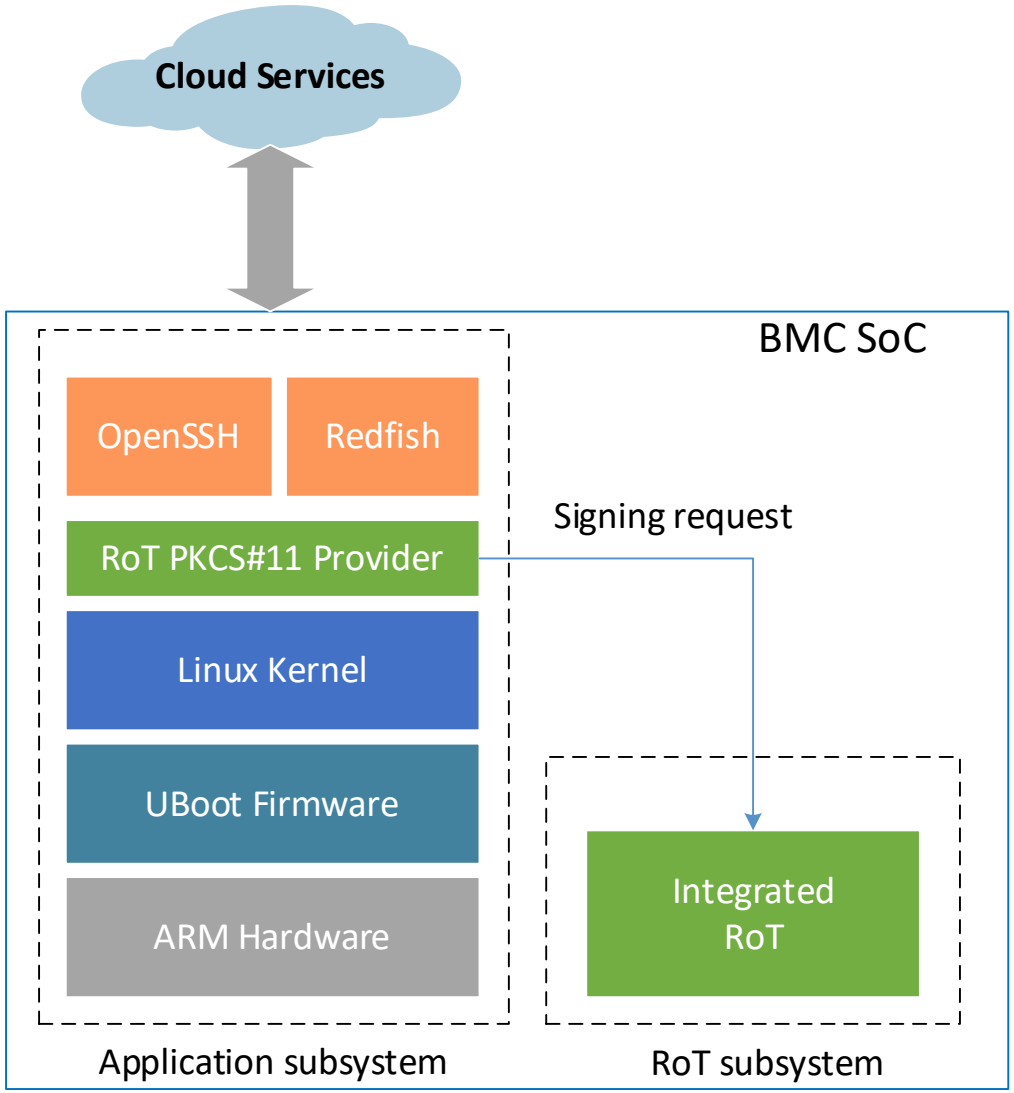

OpenSSH session establishment signing requests are directed to IROT

Microsoft

Key Insights on Enhancing Remote Access Security

- Certificate-based authentication backed by HW-RoT provides superior security for BMC management.
  -  Provides hardware isolated for key protection.
  -  Provides measurements for additional assurance.
  -  Supports lifecycle of key and certificate management.
  -  Resolves Trust On First Use issue found with SSH.
- Extendable to harden Redfish HTTPS and other protocols as well.
- Works for Linux embedded and non-Linux embedded systems.
- Scalable to more than BMCs.

Microsoft

Emerging Trends in Security: Post-Quantum Cryptography (PQC)

Designed to protect against the potential threats posed by quantum computers that breaks traditional cryptographic algorithms like RSA and ECC.

Key developments

- NIST Post-Quantum Project: focusing on standardizing PQC algorithms.

Challenges and considerations

- Performance and key size
- Compatibility with existing systems and protocols for smooth transition to PQC

On-going research & developments

- Hardware accelerator developed by Azure: Adams Bridge
- Firmware solutions for PQC public key operations.
- Standardization and industrial adoption

Microsoft

Thank you!