

# OPEN POSSIBILITIES.

## Composable Security Architectures



NOVEMBER 9-10, 2021

# Composable Security Architectures


**Andrés Lagar-Cavilla**, Principal Engineer, Google

**Alberto Muñoz**, Senior Principal Engineer, Intel

**Bryan Kelly**, Principal Engineer, Microsoft

**Indranil Banerjee**, Security Partner, Meta

**Prabhu Jayanna**, Director Product Security, AMD



**OPEN**  
PLATINUM™

OPEN POSSIBILITIES.



# Premise and Problem Statement



SECURITY

## Project Cerberus

[Cerberus](#)

**TRUSTED<sup>®</sup>  
COMPUTING  
GROUP**

[CyRes](#)



[OpenTitan](#)



**OPEN**

Compute Project

[Attestation](#)

[Secure Boot](#)

OPEN POSSIBILITIES.



[SPDM](#)



[PFR](#)



# Premise and Problem Statement

1. **Server Boot Integrity:** A system boots the intended mutable code.
2. **Remote Attestation:** A system is able to provide cryptographic proof of its identity and firmware integrity.



SECURITY

## System Integrators

“we don’t know how to build machines that will satisfy **all** customers/hyperscalers”

## Hyperscalers

*“There are too many different solutions for the same problem. OEM/OTS/OCF products do not fit our needs (yet)”*

## Suppliers

“there are so many [OCF] specs”

All: We would benefit from **alignment and consistency** to attain boot integrity and attestation

Presentation Goal: **Admissible Architectures to drive alignment & consistency**

OPEN POSSIBILITIES.



# Admissible Architecture: Approach



SECURITY

“Journey of Discovery”: we will build the idea of an admissible architecture by breaking it down into three steps

## Capability


- Establish definition of “critical devices” of a server platform
- Set “Minimum Elements” needed for critical devices

## Interoperability

- Physical interfaces (inband v/s sideband) for critical devices
- Functional APIs for “Minimum Elements” for each device

## Orchestration

- Establish a common mission for an orchestrating device (“pRoT”)
- Root critical device in a server platform hierarchy / topology



OPEN POSSIBILITIES.

API for admissible architectures

Physical Interface (Sideband)



# Server Critical Devices: Criteria




SECURITY


NIST [SP 800 193](#) definition of a “Critical Platform Device”:

“the set of platform devices necessary to minimally restore operation of the system, and sufficient to restore reasonable functionality, should themselves be resilient. We call this set of devices *critical platform devices*”


**Start with CPUs and Peripherals Handling User Data for DC**



NIC



ACCELERATOR



CPU



Storage

OPEN POSSIBILITIES.






# Server Critical Device RoT Structure

## [NIST 800-193, Platform Firmware Resiliency Guidelines](#)



SECURITY

<b>Detection</b>	<b>RTM</b> : RoT for Measurement (a.k.a. RTD)	<b>Integrity</b> TCB -- without this, can't tell what code is handling user data	<b>Integrated Silicon RoT</b> <ul style="list-style-type: none"><li>• Higher Resiliency Bar</li><li>• Protects against MITM attacks</li><li>• Limited fuses</li></ul> 
<b>Protection</b>	<b>RTU</b> : RoT for Update	<b>Availability</b> TCB -- without this, service may be denied at scale	<b>Discrete RoT Chip</b> <ul style="list-style-type: none"><li>• Platform Owner enforces Resiliency Policy<ul style="list-style-type: none"><li>• OOB updates, to what version, signed by whom</li></ul></li><li>• Integrated flash for unlimited renewability</li><li>• Single portable implementation</li></ul> 
<b>Recovery</b>	<b>RTRec</b>		

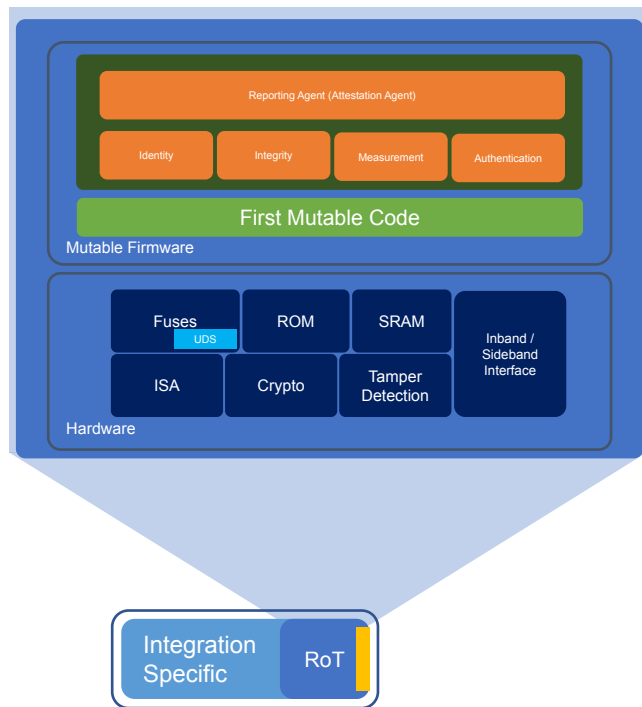


OPEN POSSIBILITIES.

# Server Critical Devices: RoT of Measurement



SECURITY



- First thing out of power-on/reset
- Must require no *programmable* power/clock external elements to function
- Measures **all** firmware for the device, without exception
- Measures (and protects) life cycle and HW state fuses
- Controls internal reset for all sub-systems in the device
- Provides (and protects) a unique device identity
- Provides unforgeable attestation rooted in unique identity

OPEN POSSIBILITIES.






# RTM Identity

RTM firmware wields **identity certificate chain**, rooted in **device hardware**

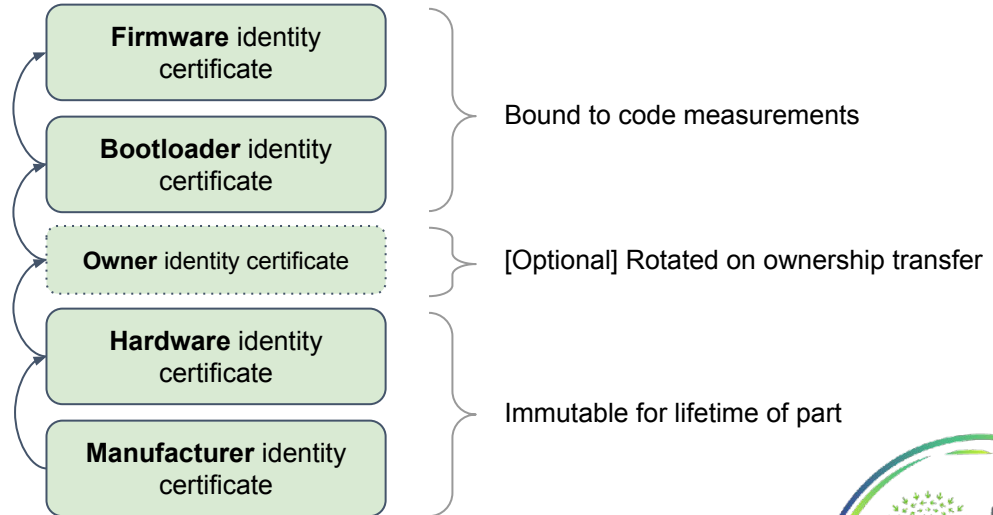


SECURITY

Each boot layer measures the next



Each identity certifies the next



Mutable RTM code must use secure boot:  
No unsigned code runs within the RTM perimeter

OPEN POSSIBILITIES.





# RTU - Update and Protection

For large, automated fleets, the **update** path is a **DoS at scale risk**




SECURITY

Untrusted  
Out of Band



ASIC



Flash

OPEN POSSIBILITIES.

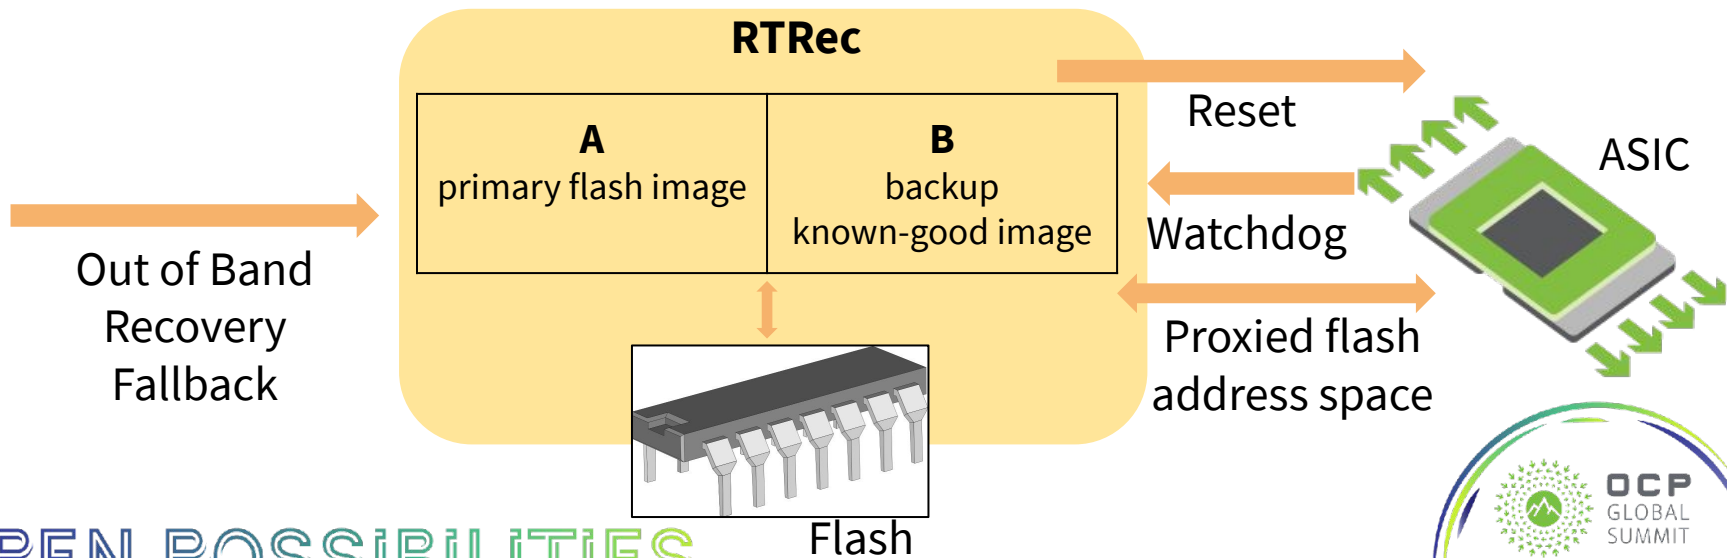
# RTRec - Recovery



SECURITY

For large, automated fleets, operator-based **repairs** are **not scalable**

- Need an additional layer of automated protection for availability
- Risk: Adversaries and bugs

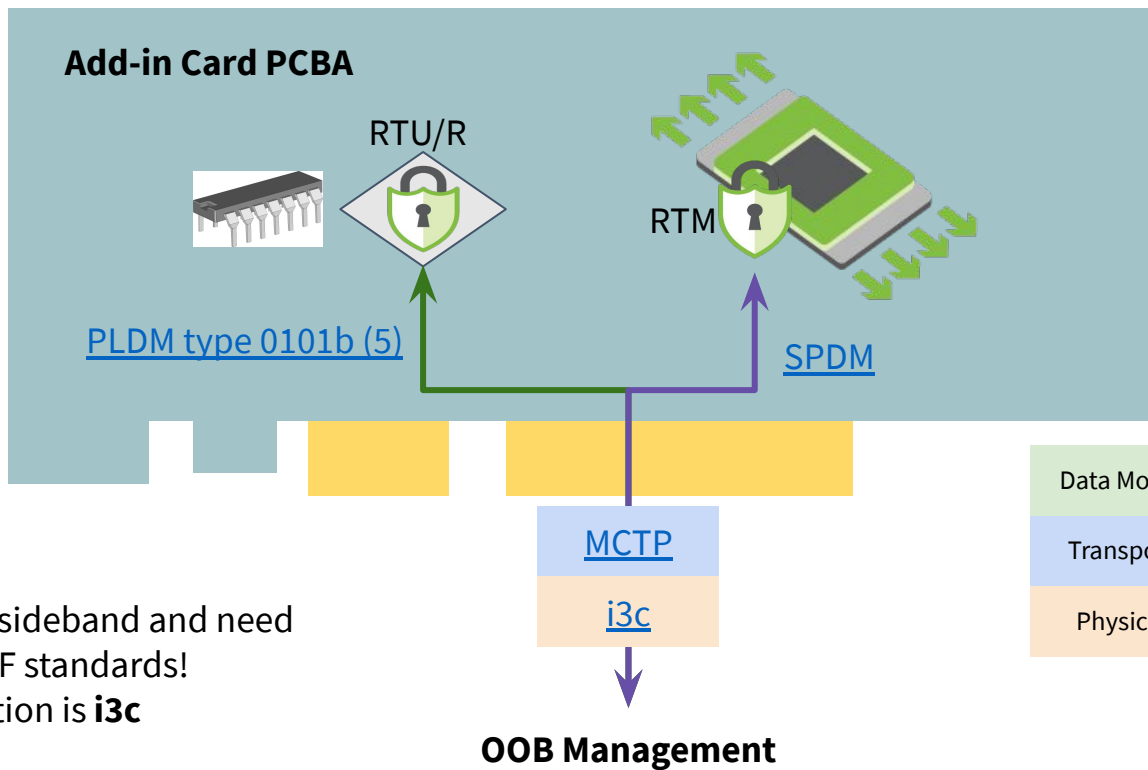


OPEN POSSIBILITIES.

# Standard Interface Alignment



SECURITY



We lack a good sideband and need  
PCI-SIG, SSD SFF standards!  
Proposed direction is **i3c**

OPEN POSSIBILITIES.

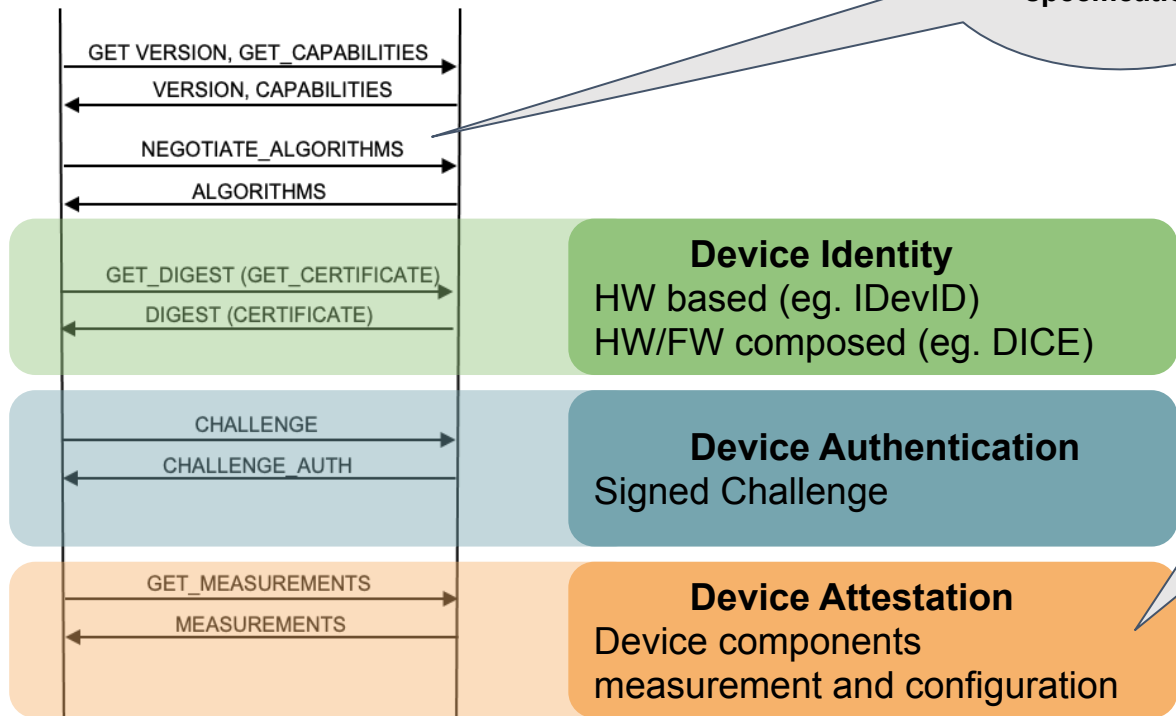
	RTU/R FW update	RTM Attestation
Data Model	<a href="#">PLDM t5</a>	<a href="#">SPDM</a>
Transport	<a href="#">MCTP</a>	
Physical	<a href="#">i3c</a>	



# SPDM Profile

Security Orchestrator

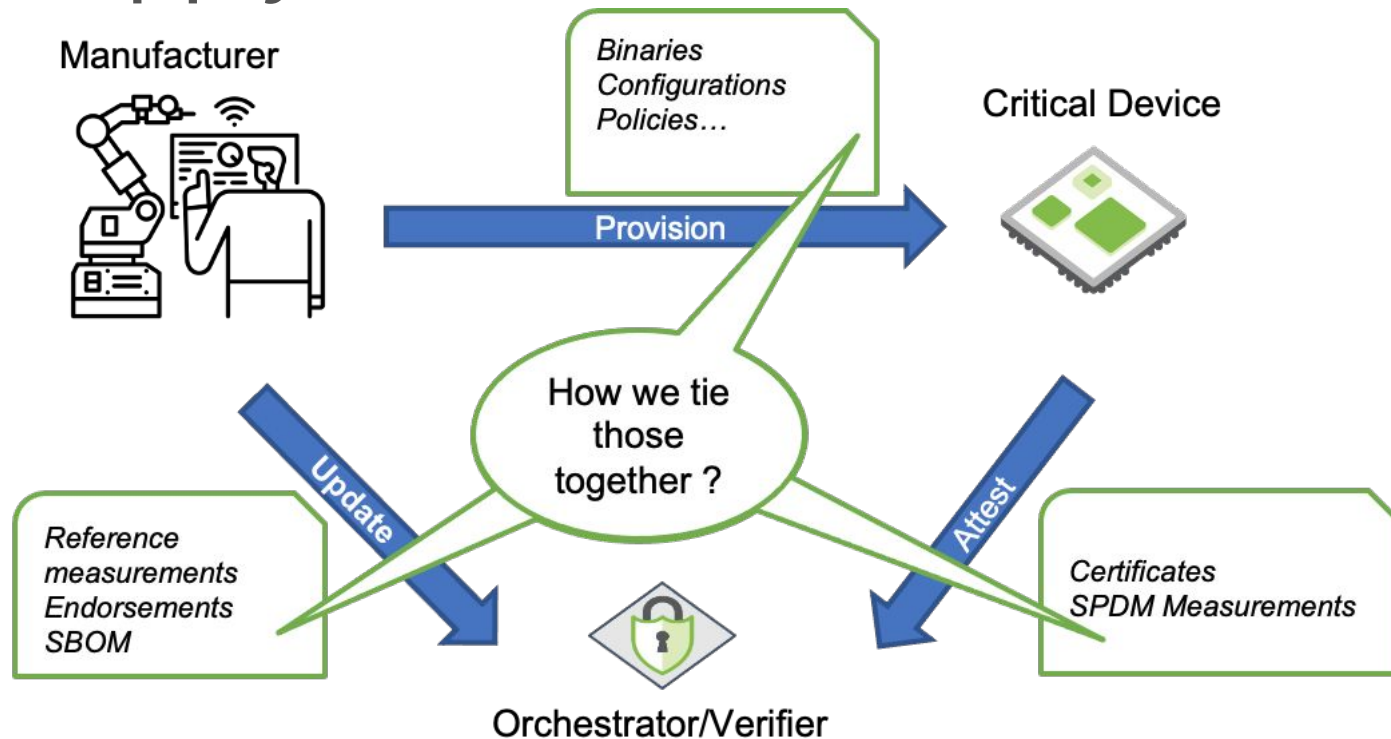
Critical Device



SECURITY

OPEN POSSIBILITIES.

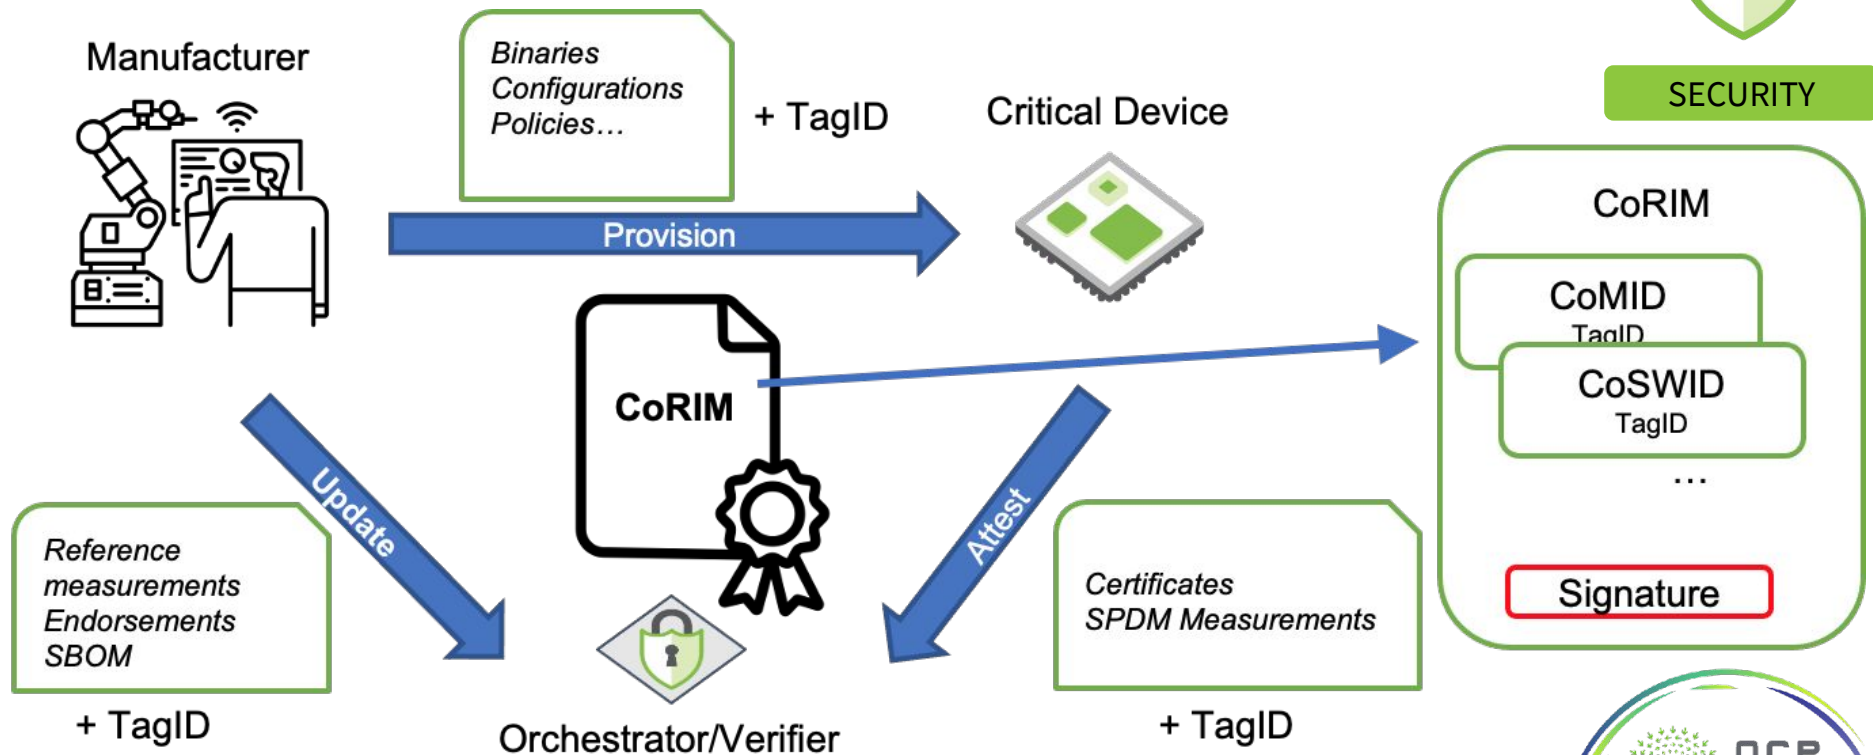
# Supply Chain of Measurements



SECURITY

OPEN POSSIBILITIES.

# Concise Reference Integrity Manifest (CoRIM)



OPEN POSSIBILITIES.

# Call to Action (1/2)



SECURITY

- **Define** standardized **FW descriptor** and signing
  - TCG CO-RIM going to ballot in 2021, aiming for 2022/2023 1.0
- OCP **defined** manifest format for **SPDM reporting** of attestation **blocks**
  - Using CoMID, CoSWID and Tag IDs – join Piotr and Alberto in Security WG calls
- **Need MCTP hw channel** in all phys connector specifications
  - Not i2c. We propose **i3c**
  - **Requirements:** Out of band, scalable, switchable, bi-directional, high bandwidth (ex: push 1Gbit of fw in seconds)
  - HW Specifications: PCI CEM @PCI-SIG, small form factor SSD, OCP NIC 3.0

OPEN POSSIBILITIES.





# An Admissible Arch Coming Together

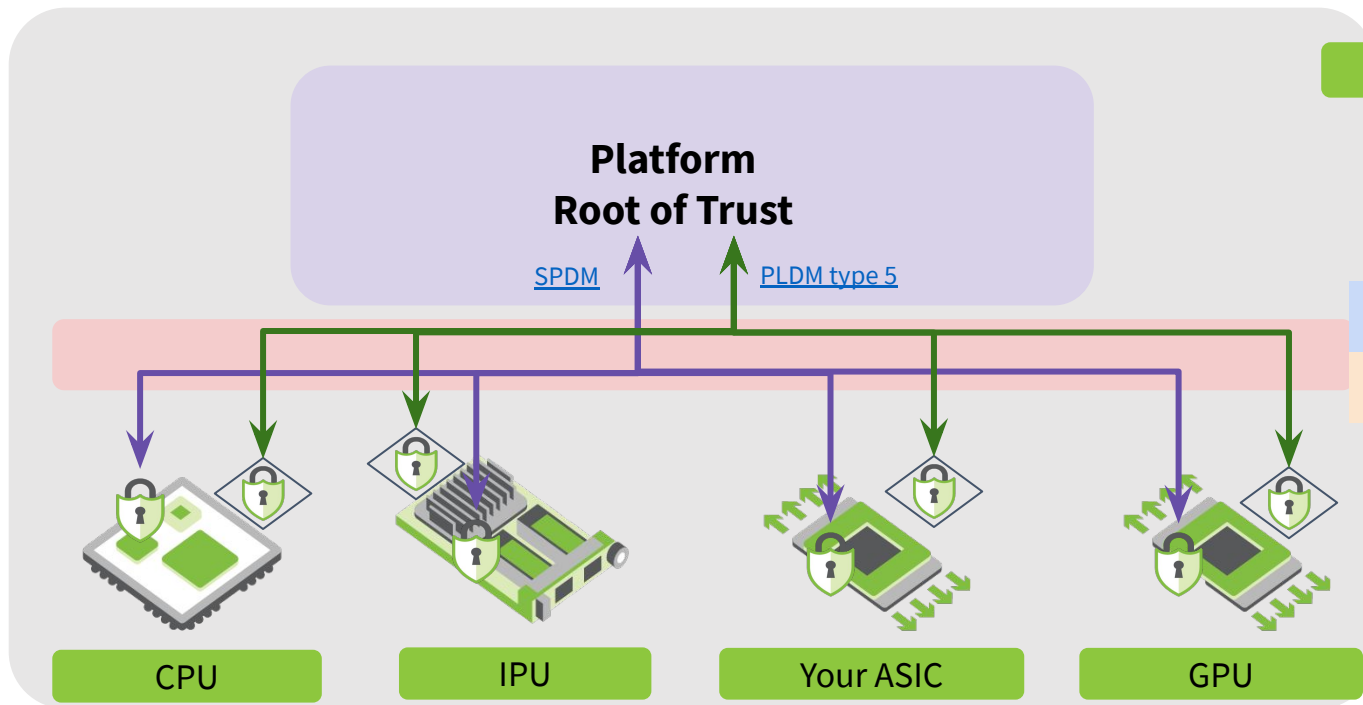


SECURITY

Single  
OOB bus

MCTP

i3c



OPEN POSSIBILITIES.

# Admissible Arch Platform Root of Trust



## [NIST 800-193, Platform Firmware Resiliency Guidelines](#)

SECURITY

<b>Detection</b>	<b>Aggregate</b> all RTM measurements, <b>verify fitness</b> of the system against <b>policy</b> of target firmware	<ul style="list-style-type: none"><li>• <b>Detection:</b> Control Plane aggregates measurements from all RTMs<ul style="list-style-type: none"><li>○ <b>Trusted fabric service</b> verifies fitness and authorizes service</li></ul></li></ul>
<b>Protection</b>	Centralize <b>out of band update</b> flow	<ul style="list-style-type: none"><li>• BMC centralizes OOB fw updating</li></ul>
<b>Recovery</b>	Provide out of band recovery with <b>golden images</b> to RTRec of unresponsive critical devices	<ul style="list-style-type: none"><li>• <b>pRoT</b> chip is the <b>minimum recovery foothold</b><ul style="list-style-type: none"><li>○ pRoT is its own RTM/RTU/RTRec</li><li>○ Bootstraps recovery of BMC</li><li>○ BMC supplies recovery images for critical devices</li></ul></li></ul>



OPEN POSSIBILITIES.



# Composable Admissible Architecture



SECURITY

Single  
OOB bus

MCTP

i3c



Trusted Verifier

DC-SCM Modularity

Platform Root of Trust

BMC



pRoT

SPDM

PLDM type 5

Internal RTM  
for attestation



Discrete RTU +  
RTRec

PLDM fw  
upgrade  
OOB path

SPDM  
attestation  
path



CPU



IPU



Your ASIC



GPU

OPEN POSSIBILITIES.

OCP Platform  
Security Overview



OCP  
GLOBAL  
SUMMIT

NOVEMBER 9-10, 2021

# Composable Admissible Architecture



SECURITY

Tl;dr list for offline viewing -- alignment on

- [NIST 800-193](#) structure with RTM in SoC, discrete RTU/RTRec in PCB
- [DICE](#)-like key derivation for RTM renewable security
- [SPDM](#) 1.2+ as attestation lingua franca
  - Key derivation and transport consistent with [OCP Attestation](#)
  - Call to action on Co-RIM for fw descriptor, CoMID/CoSWID for SPDM attestation blocks
- OOB fw push path: [PLDM](#)
- Call to action on scalable, bi-directional, out of band path for [MCTP](#): [i3c](#)
- Use of OCP [DC-SCM](#) for system and PRoT modularity
- Disaggregated PRoT model with trusted external verifier
- Documented in progress at [OCP Platform Security Overview](#)
  - Consistent with [OCP threat model](#) and [security checklist](#)

## OPEN POSSIBILITIES.



# Call to Action (2/2)

- **Contribute** to [OCP Platform Root of Trust](#) document
- Make **MCTP** scalable, switchable, high-bandwidth i3c **sidebands** a standard
  - HW Specifications: PCI CEM @PCI-SIG, small form factor SSD, OCP NIC 3.0
- **Define** standardized **FW descriptor** and signing
  - TCG CO-RIM going to ballot in 2021, aiming for 2022/2023 1.0
- OCP **defined** manifest format for **SPDM reporting** of attestation **blocks**
  - Using CoMID, CoSWID and Tag IDs – join Piotr & Alberto in Security WG calls
- Timeline for **OCP docs ratification**
  - [Attestation](#), [Checklist](#), [Secure Boot](#), [Platform Security Overview](#)
- Upcoming silicon products that comply as “critical devices”

Project Wiki with latest specification : <https://www.opencompute.org/wiki/Security>

OPEN POSSIBILITIES.



# Open Discussion



NOVEMBER 9-10, 2021

# BACKUP



NOVEMBER 9-10, 2021

# Critical Devices in Server: Minimal Elements

Resiliency RoT Service	Critical Device	Orchestrator Device
RoT for Detection (RTD)	Mandatory	Mandatory
RoT for Update (RTU)	Optional	Mandatory
RoT for Recovery (RTRec)	Optional	Mandatory

## Minimum Elements to Enable RTD

- Identity
- Confidentiality
- Integrity
- Authentication
- Measurement

OPEN POSSIBILITIES.




# Alternate View Critical Device RoT



## [NIST 800-193, Platform Firmware Resiliency Guidelines](#)

SECURITY

Detection	<b>RTM</b> (RoT for Measurement) a.k.a. RTD	<b>Integrity</b> TCB -- without this, can't tell what code is handling user data
Protection	<b>RTU</b> (RoT for Update)	<b>Availability</b> TCB -- without this, service may be denied at scale
Recovery	<b>RTRec</b>	




	Discrete RoT	Integrated RoT
Strengths	<ul style="list-style-type: none"><li>Platform level Resiliency Policy Enforcement</li><li>Enables Platform Owner to be in control</li><li>Integrated flash for unlimited renewability</li></ul>	<ul style="list-style-type: none"><li>Higher Resiliency Bar</li><li>Resiliency Policy within limited fuses</li><li>Self protects against MITM attacks (some)</li></ul>
Weaknesses	<ul style="list-style-type: none"><li>MITM susceptibility while interacting with other devices</li></ul>	<ul style="list-style-type: none"><li>No integrated flash</li><li>Fragmented resiliency policies and interoperability challenge</li></ul>

OPEN POSSIBILITIES.

# RoT Structure Alternatives

## Discrete RoT: All functions




Titan, Cerberus, CEC1712, ...

**Strength:** node process affords flash. Unlimited rotations, rollback control

**Weakness:** SPI interposition defeats RTM

## Combo: Discrete + Internal



**Strength:** Single RTU + RTRec implementation for all devices

Best of both worlds

**Weakness:** extra BOM, space challenge on some cards

## Internal RoT: all fns



**Strength:** Defeats SPI interposition, secure boot channel

**Weakness:** Limited OTP bank for rotations, rollback control. RTU/RTRec complexity across suppliers

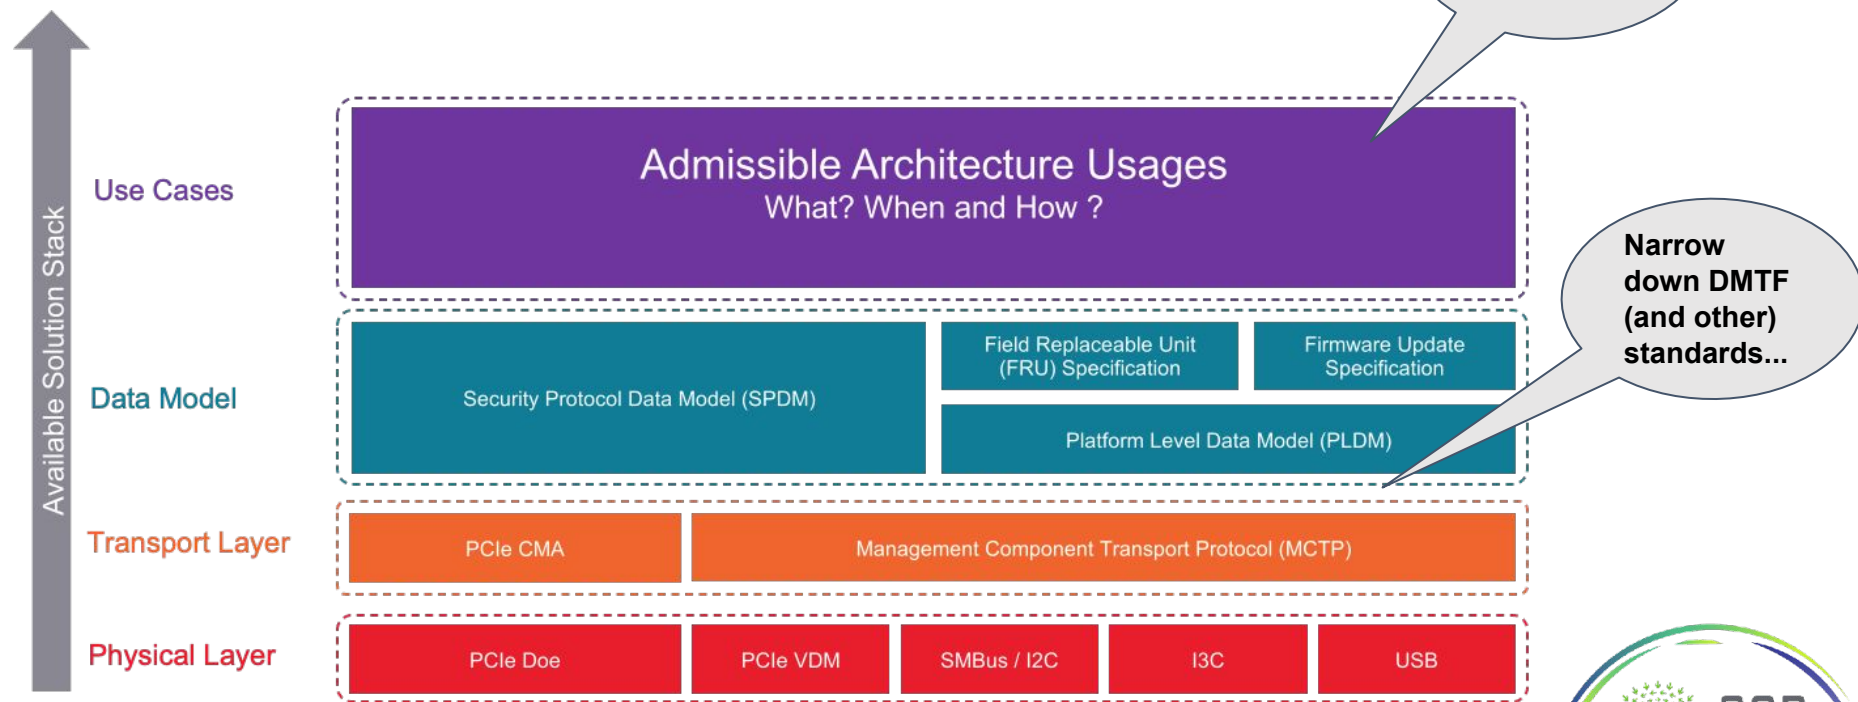


SECURITY

**Preferred**

OPEN POSSIBILITIES.

# Standard Protocol Layers



OPEN POSSIBILITIES.


# Renewable DICE problem

**Immutable keys** should be wielded by **immutable code**.  
But what if there's a bug in the code that wields that key?

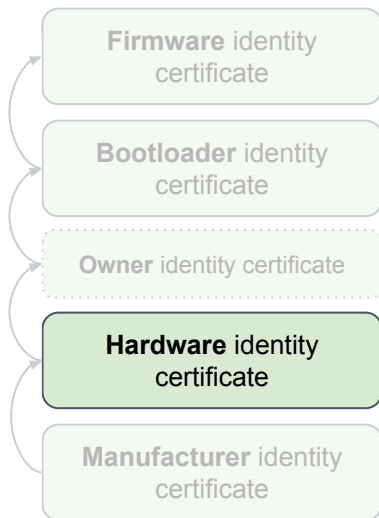


SECURITY

Each boot layer measures the next



Mutable RTM code must use secure boot:  
No unsigned code runs within the RTM perimeter



Options:

1. Hardware identity is wielded by mask ROM.
2. Hardware identity is wielded by FMC (first-mutable-code) and is bound to FMC's code identity.
  - a. FMC is effectively immutable.
3. Hardware identity is wielded by FMC and is bound to FMC's *signer identity*.
  - a. FMC is mutable, but its code identity cannot be attested.

Available for  
**nascent** implementations

Recommended for  
**mature** implementations

OPEN POSSIBILITIES.



# SPDM Profile - measurement manifest



SECURITY

- FW Manifest and signing
  - SBOM produced by manufacturers
  - SWID – unified Software Identification Tag – xml based
  - CoSWID – Concise SoftWare Identification Tag – JSON based
  - CoMID – Concise Module Identification Tag. For HW and embedded FW – CDDL/JSON based
  - CoRIM – Concise Reference Integrity Manifest. Signing container for CoMID and CoSWID
- Measurement manifest:
  - Use TCG CoMID as baseline
  - Direct mapping between CoMID and SPDM Measurement Manifest (i.e., provide CoMID schema as additional format for Measurement Manifest)
- Define minimum set of measurements required per device class

OPEN POSSIBILITIES.



# SPDM Measurement - per device class



SECURITY

- Different class of devices might need different minimum set of measurements required. Some examples:

## CPU SoC

### Measurements:

- all internal uC FW components
- HW Security Configuration
- Fuses
- x86/ARM cores boot code
- ...

## SSD Drive

### Measurements:

- main FW
- on board uC FW
- ...

## FPGA

### Measurements:

- FPGA bitstream
- on board uC FW
- ...

OPEN POSSIBILITIES.




# An Admissible Arch Coming Together



SECURITY

 Internal  
RTM

 Discrete RTU  
+ RTRec



Intel IPU



Google TPU



Microsoft Accel



AMD GPU



Your ASIC



Intel Xeon CPU



AMD Zen CPU

OPEN POSSIBILITIES.

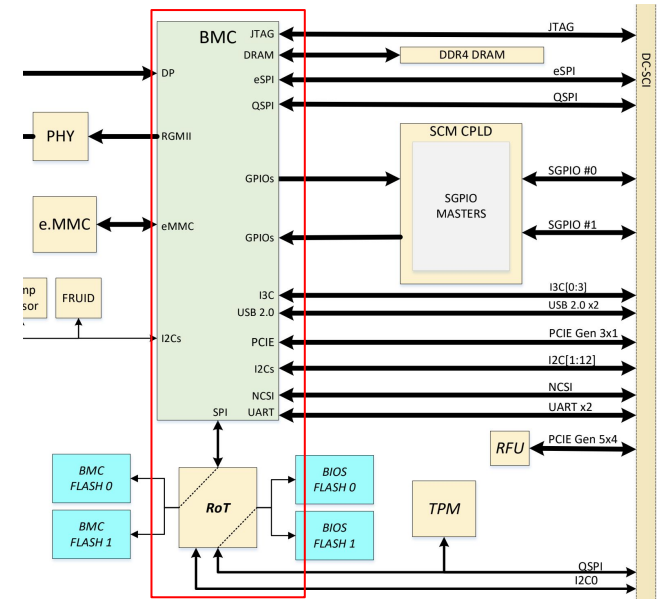


# PRoT and DC-SCM

- OCP [DC-SCM](#) provides a natural “house” for PRoT
- RTM in CPU neutralizes DC-SCI interposition attack
- BMC within SCM behaves as availability only actor
  - BMC can route SPDM/PLDM traffic, attestation challenges and fw images
  - End-to-end crypto keeps BMC out of boot TCB
  - BMC can always shut down the system -- but cannot brick with bad fw or forge attestations
- SCM enables PRoT modularity



SECURITY



OPEN POSSIBILITIES.




# pRoT Alternatives Considered



SECURITY

## Preferred: Disaggregated




- Cerberus, PFR, Titan: chip is RTM/U/Rec for BMC
- RoT chip is the minimum recovery foothold
  - Bootstraps recovery of an entire system
  - BMC media can supply recovery images for leafs
- BMC centralizes OOB fw updating
- BMC aggregates measurements from all RTMs
- Trusted fabric service verifies fitness and authorizes service

## Dedicated Chip



- Chip verifies fitness against manifest deployed to its NVM

## All-in-one

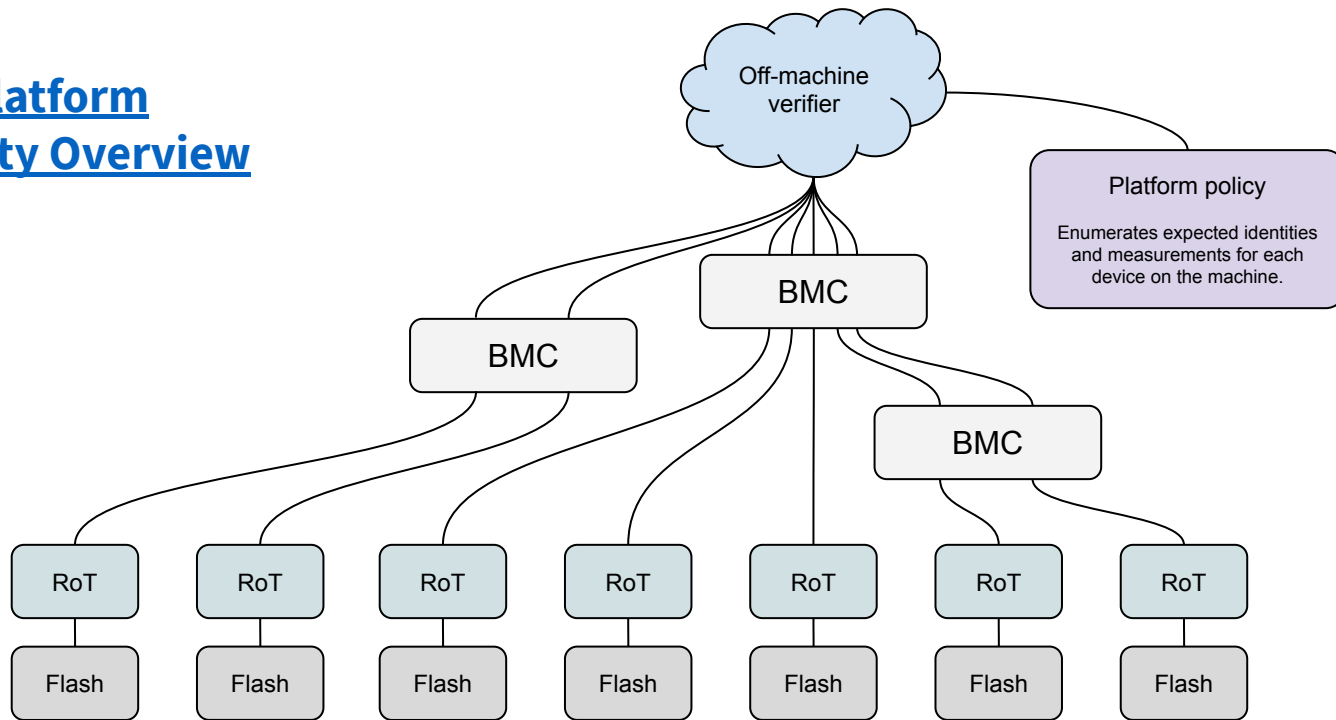


- BMC + RoT are a single PProT ASIC.
- All functions merged: BMC integrity, system recovery, fitness
- Typical in OEMs
- May not stay in preso

OPEN POSSIBILITIES.

# Disaggregated PRoT Drawing

## OCP Platform Security Overview



OPEN POSSIBILITIES.