BEDROCK
Systems Inc

Using the NOVA Microhypervisor for Trusted Computing at Scale

Udo Steinberg

Agenda

- NOVA Microhypervisor Overview
  - Architecture, Scalability
  - Innovation Timeline
- Trusted Computing
  - Verified vs. Measured Boot / SRTM vs. DRTM
  - Intel Trusted Execution Technology (TXT)
  - Trusted Platform Module (TPM)
  - Integrity Measurement of NOVA and Root-PD
- Q & A

BedRock Ultravisor Architecture

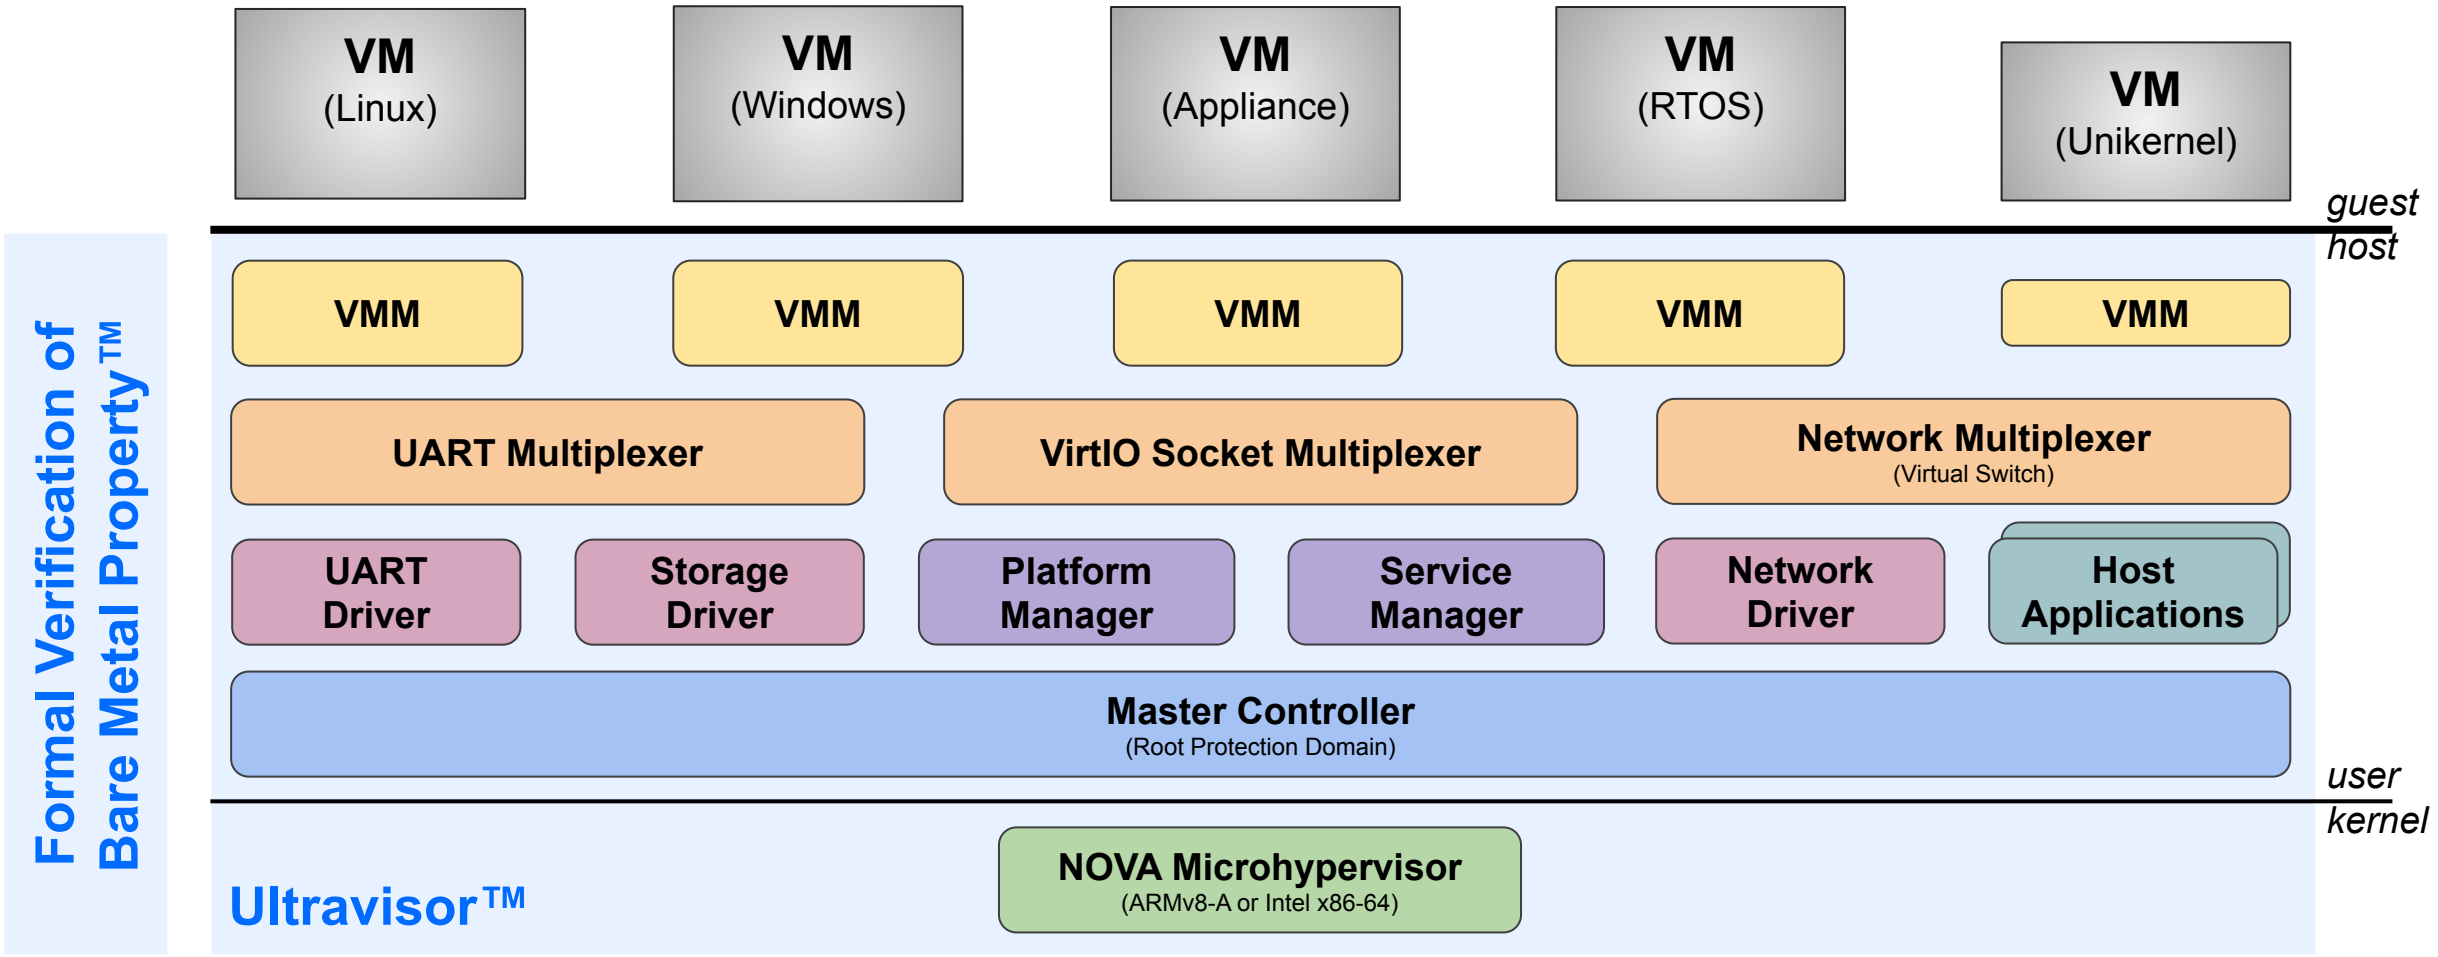

Scaling NOVA from Embedded to Cloud Servers

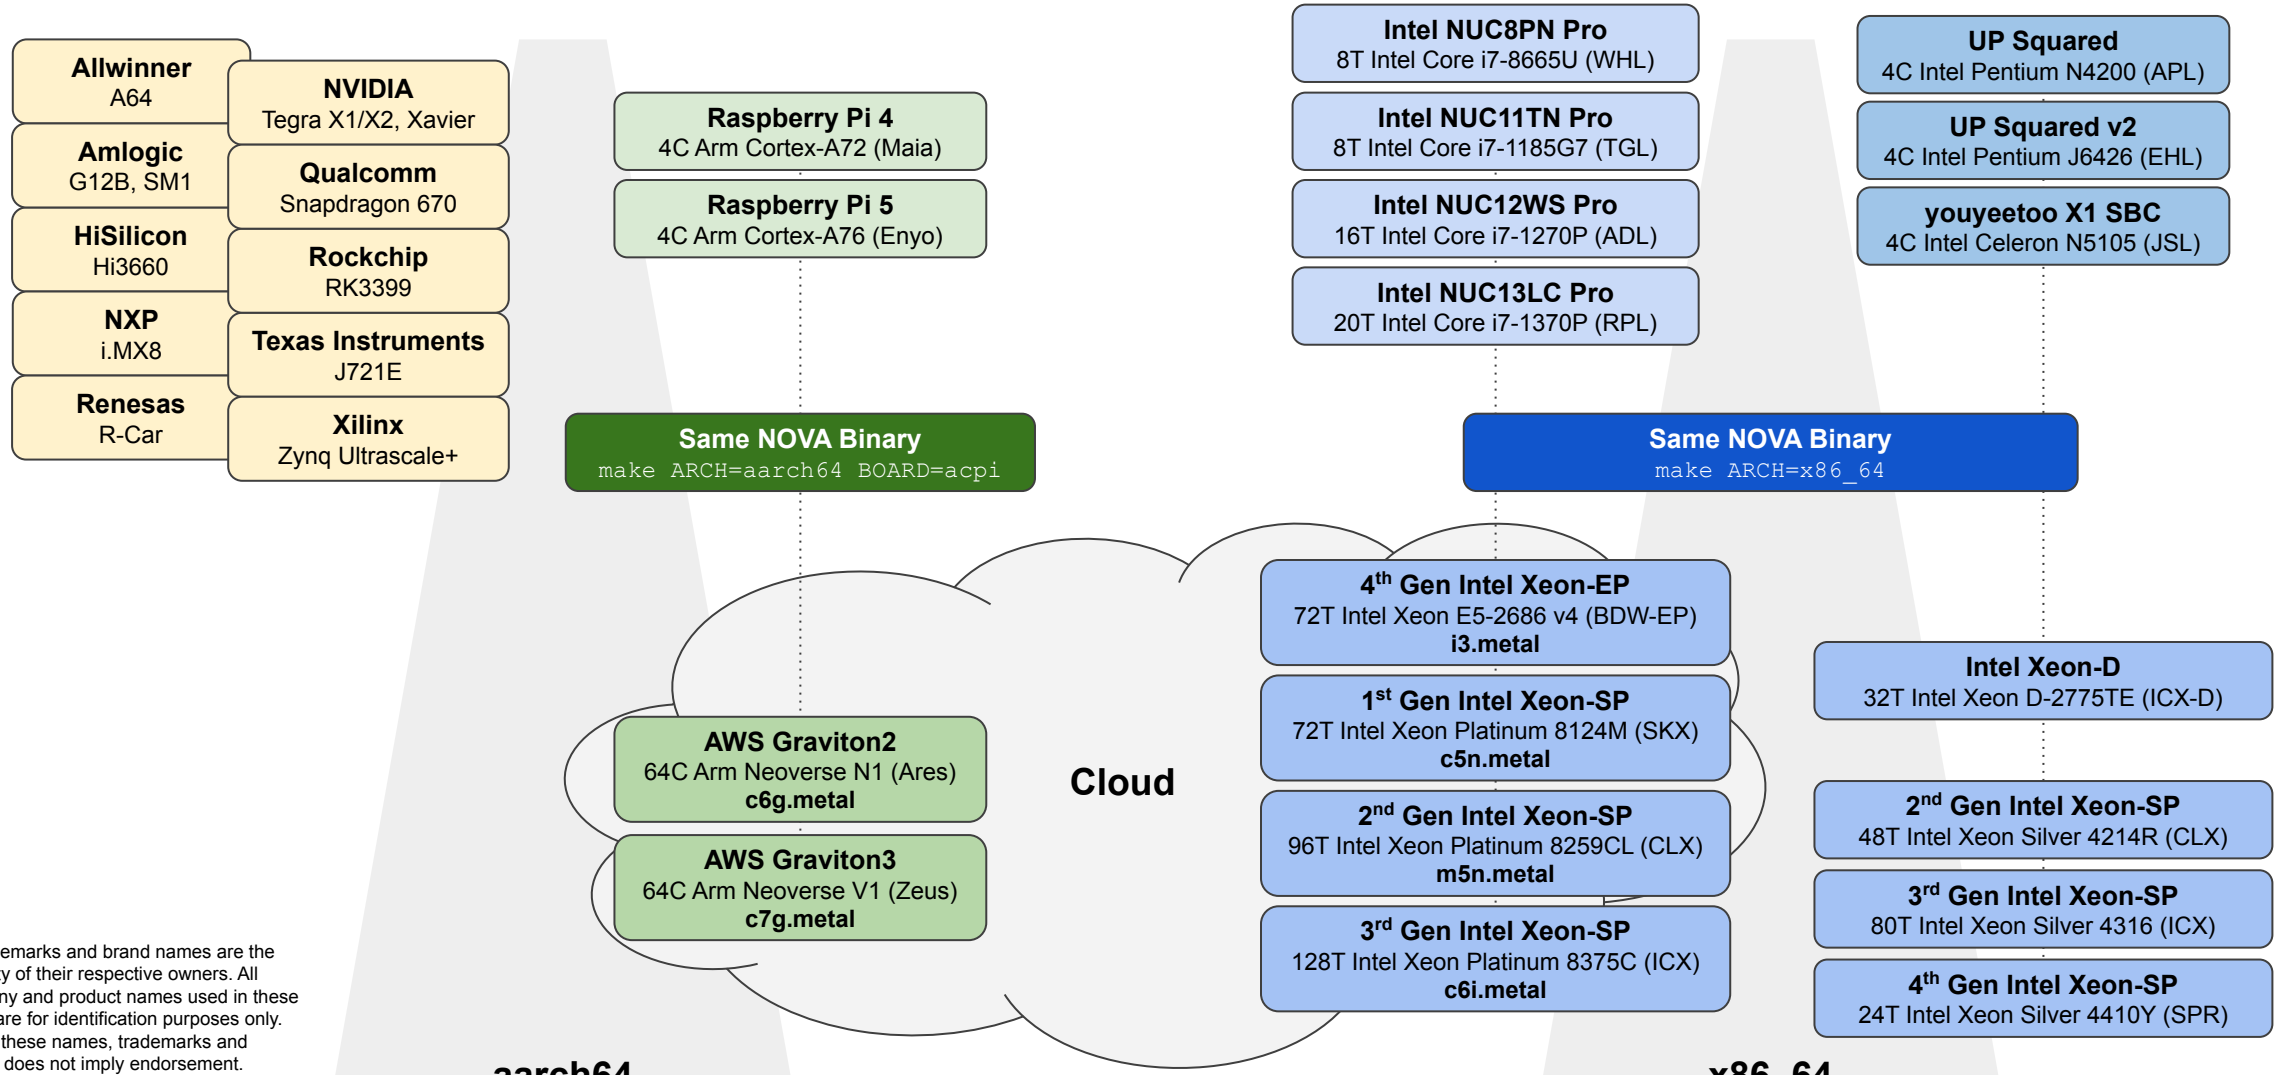

NOVA / Embedded: Raspberry Pi 4 and 5

```
NOVA Microhypervisor #8b6c28a (aarch64): Feb  1 2024 10:57:11 [gcc 13.2.0]

[ -1] XSDT: 0x33b2fe98 OEM:RPIFDN TBL:RPI4     REV: 1 LEN:    116 (0xb8) ok
[ -1] FACP: 0x33b2e998 OEM:RPIFDN TBL:RPI4     REV: 6 LEN:    276 (0x3f) ok
[ -1] CSRT: 0x33b2fa98 OEM:RPIFDN TBL:RPI4     REV: 0 LEN:    361 (0x74) ok
[ -1] DBG2: 0x33b2fd18 OEM:RPIFDN TBL:RPI4     REV: 0 LEN:     97 (0x97) ok
[ -1] GTDT: 0x33b2f998 OEM:RPIFDN TBL:RPI4     REV: 3 LEN:    104 (0xea) ok
[ -1] IORT: 0x33b2f598 OEM:RPIFDN TBL:RPI4     REV: 0 LEN:    183 (0x00) ok
[ -1] APIC: 0x33b2f698 OEM:RPIFDN TBL:RPI4     REV: 5 LEN:    388 (0x5f) ok
[ -1] PPTT: 0x33b2eb18 OEM:RPIFDN TBL:RPI4     REV: 2 LEN:    388 (0xfa) ok
[ -1] SPCR: 0x33b2f918 OEM:RPIFDN TBL:RPI4     REV: 2 LEN:     80 (0x0e) ok
[ -1] SSDT: 0x33b2ed98 OEM:RPIFDN TBL:RPI4EMMC REV: 2 LEN:    631 (0x29) ok
[ -1] SSDT: 0x33b2a198 OEM:RPIFDN TBL:RPI4PCIE REV: 5 LEN:    660 (0xbc) ok
[ -1] ACPI: Version 6.3 Profile 6 Features 0x100021 Boot 0x1
[ -1] PSCI: Version 1.1 States 0xb0
[ -1] SPCR: Console 8000:0003 (0:0xfe201000:32:3)
[ -1] DBG2: Console 8000:0003 (0:0xfe201000:32:3)
[  0] CORE: 00:00:00:00 Arm Cortex-A72 (Maia) r0p3 PA:4 XNX:0 GIC:0 (EL2)
[  0] GICD: 0xff841000 v2 r0p1 Impl:0x43b Prod:0x2 ESPI:0 LPIS:0 INT:256 S:1 G:1
[  0] GICC: 0xff842000 v2 r2p1 Impl:0x43b Prod:0x20
[  0] GICH: 0xff844000 APR:1 LR:4
[  0] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  1] CORE: 00:00:00:01 Arm Cortex-A72 (Maia) r0p3 PA:4 XNX:0 GIC:0 (EL2)
[  1] GICC: 0xff842000 v2 r2p1 Impl:0x43b Prod:0x20
[  1] GICH: 0xff844000 APR:1 LR:4
[  1] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  2] CORE: 00:00:00:02 Arm Cortex-A72 (Maia) r0p3 PA:4 XNX:0 GIC:0 (EL2)
[  2] GICC: 0xff842000 v2 r2p1 Impl:0x43b Prod:0x20
[  2] GICH: 0xff844000 APR:1 LR:4
[  2] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  3] CORE: 00:00:00:03 Arm Cortex-A72 (Maia) r0p3 PA:4 XNX:0 GIC:0 (EL2)
[  3] GICC: 0xff842000 v2 r2p1 Impl:0x43b Prod:0x20
[  3] GICH: 0xff844000 APR:1 LR:4
[  3] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  0] TIME: 371ms 150ms/142ms
[  0] ROOT: No image
```

```
NOVA Microhypervisor #8b6c28a (aarch64): Feb  1 2024 10:57:11 [gcc 13.2.0]

[ -1] XSDT: 0x3867fe98 OEM:RPIFDN TBL:RPI5     REV: 1 LEN:     84 (0x56) ok
[ -1] FACP: 0x3867fb98 OEM:RPIFDN TBL:RPI5     REV: 6 LEN:    276 (0xdc) ok
[ -1] DBG2: 0x3867fa98 OEM:RPIFDN TBL:RPI5     REV: 0 LEN:     97 (0x1f) ok
[ -1] GTDT: 0x3867fd18 OEM:RPIFDN TBL:RPI5     REV: 3 LEN:    104 (0x8c) ok
[ -1] APIC: 0x3867e998 OEM:RPIFDN TBL:RPI5     REV: 5 LEN:    388 (0x53) ok
[ -1] PPTT: 0x3867f698 OEM:RPIFDN TBL:RPI5     REV: 2 LEN:    304 (0xdf) ok
[ -1] SPCR: 0x3867fe18 OEM:RPIFDN TBL:RPI5     REV: 2 LEN:     80 (0x9e) ok
[ -1] ACPI: Version 6.3 Profile 6 Features 0x100021 Boot 0x1
[ -1] PSCI: Version 1.1 States 0xb0
[ -1] SPCR: Console 8000:0003 (0:0x107d001000:32:3)
[ -1] DBG2: Console 8000:0003 (0:0x107d001000:32:3)
[  0] CORE: 00:00:00:00 Arm Cortex-A76 (Enyo) r4p1 PA:2 XNX:1 GIC:0 (EL2)
[  0] GICD: 0x107fff9000 v2 r0p1 Impl:0x43b Prod:0x2 ESPI:0 LPIS:0 INT:320 S:1 G:1
[  0] GICC: 0x107fffa000 v2 r2p1 Impl:0x43b Prod:0x20
[  0] GICH: 0x107fffc000 APR:1 LR:4
[  0] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  1] CORE: 00:00:01:00 Arm Cortex-A76 (Enyo) r4p1 PA:2 XNX:1 GIC:0 (EL2)
[  1] GICC: 0x107fffa000 v2 r2p1 Impl:0x43b Prod:0x20
[  1] GICH: 0x107fffc000 APR:1 LR:4
[  1] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  2] CORE: 00:00:02:00 Arm Cortex-A76 (Enyo) r4p1 PA:2 XNX:1 GIC:0 (EL2)
[  2] GICC: 0x107fffa000 v2 r2p1 Impl:0x43b Prod:0x20
[  2] GICH: 0x107fffc000 APR:1 LR:4
[  2] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  3] CORE: 00:00:03:00 Arm Cortex-A76 (Enyo) r4p1 PA:2 XNX:1 GIC:0 (EL2)
[  3] GICC: 0x107fffa000 v2 r2p1 Impl:0x43b Prod:0x20
[  3] GICH: 0x107fffc000 APR:1 LR:4
[  3] TIMR: EL2p:10L EL1v:11L 54000000 Hz
[  0] TIME: 185ms 0ms/105ms
[  0] ROOT: No image
```

Both using UEFI firmware with ACPI support

NOVA / Server: Arm Neoverse V1 (c7g.metal)

```
NOVA Microhypervisor #8b6c28a (aarch64): Feb  1 2024 10:57:11 [gcc 13.2.0]

[ -1] XSDT: 0x15e1fe98 OEM:AMAZON TBL:GRVTN003 REV: 1 LEN:    140 (0x9f) ok
[ -1] FACP: 0x15e1fb98 OEM:AMAZON TBL:GRVTN003 REV: 6 LEN:    276 (0x30) ok
[ -1] HEST: 0x15e1fa98 OEM:AMAZON TBL:GRVTN002 REV: 1 LEN:    168 (0x8a) ok
[ -1] GTDT: 0x15e1fd18 OEM:AMAZON TBL:GRVTN003 REV: 2 LEN:     96 (0xbb) ok
[ -1] APIC: 0x15e17518 OEM:AMAZON TBL:GRVTN003 REV: 5 LEN:   5268 (0x45) ok
[ -1] SRAT: 0x15e1d898 OEM:AMAZON TBL:GRVTN003 REV: 3 LEN:   1496 (0x2f) ok
[ -1] SLIT: 0x15e1fe18 OEM:AMAZON TBL:GRVTN003 REV: 1 LEN:     45 (0x24) ok
[ -1] MCFG: 0x15e1df18 OEM:AMAZON TBL:GRVTN003 REV: 1 LEN:    108 (0xad) ok
[ -1] PPTT: 0x15e19098 OEM:AMAZON TBL:GRVTN003 REV: 2 LEN:   3262 (0xe4) ok
[ -1] SDEI: 0x15e1e918 OEM:AMAZON TBL:GRVTN003 REV: 1 LEN:     36 (0x4f) ok
[ -1] IORT: 0x15e1e018 OEM:AMAZON TBL:GRVTN003 REV: 0 LEN:    800 (0xb2) ok
[ -1] SSDT: 0x15e1ce18 OEM:AMAZON TBL:GRVTN003 REV: 2 LEN:   1500 (0x8c) ok
[ -1] SSDT: 0x15e1e518 OEM:AMAZON TBL:GRVTN003 REV: 2 LEN:     74 (0xed) ok
[ -1] SPCR: 0x15e1ff98 OEM:AMAZON TBL:GRVTN003 REV: 2 LEN:     80 (0x86) ok
[ -1] ACPI: Version 6.1 Profile 4 Features 0x301000 Boot 0x1
[ -1] PSCI: Version 1.1 States 0xb0
[ -1] SPCR: Console 8000:0000 (0:0xe2f00000:8:1)
[ -1] PCIE: 0xe000000000 Segment 0x0000 Bus 0x00-0xff
[ -1] PCIE: 0xe010000000 Segment 0x0001 Bus 0x00-0xff
[ -1] PCIE: 1d0f:0200 06-04-00 D0 PCIE PMI           0001:00:00.0
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI              0001:01:00.0
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.0
[ -1] PCIE: 1d0f:8250 07-00-03 D0 PCIE PMI                    0001:03:00.0
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.1
[ -1] PCIE: 1d0f:0061 01-08-02 D0 PCIE PMI FLR                0001:04:00.0
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.2
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.3
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.4
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.5
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.6
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:00.7
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.0
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.1
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.2
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.3
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.4
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.5
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.6
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:01.7
[ -1] PCIE: 1d0f:cec2 06-04-00 D0 PCIE PMI                 0001:02:02.0
```

```
[ 56] CORE: 01:24:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 56] GICR: 0xbc00740000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 56] GICC: REGS
[ 56] GICH: REGS APR:1 LR:4
[ 56] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[ 57] CORE: 01:25:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 57] GICR: 0xbc00780000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 57] GICC: REGS
[ 57] GICH: REGS APR:1 LR:4
[ 57] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[ 58] CORE: 01:26:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 58] GICR: 0xbc007c0000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 58] GICC: REGS
[ 58] GICH: REGS APR:1 LR:4
[ 58] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[ 59] CORE: 01:27:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 59] GICR: 0xbc00800000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 59] GICC: REGS
[ 59] GICH: REGS APR:1 LR:4
[ 59] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[ 60] CORE: 01:28:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 60] GICR: 0xbc00840000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 60] GICC: REGS
[ 60] GICH: REGS APR:1 LR:4
[ 60] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[ 61] CORE: 01:29:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 61] GICR: 0xbc00880000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 61] GICC: REGS
[ 61] GICH: REGS APR:1 LR:4
[ 61] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[ 62] CORE: 01:30:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 62] GICR: 0xbc008c0000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 62] GICC: REGS
[ 62] GICH: REGS APR:1 LR:4
[ 62] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[ 63] CORE: 01:31:00:00 Arm Neoverse V1 (Zeus) r1p1 PA:5 XNX:1 GIC:3 (EL2)
[ 63] GICR: 0xbc00900000 v4 r0p1 Impl:0x43b Prod:0x4 EPPI:1 MPAM:1
[ 63] GICC: REGS
[ 63] GICH: REGS APR:1 LR:4
[ 63] TIMR: EL2p:10L EL1v:11L 1050000000 Hz
[  0] TIME: 8842ms 16ms/786ms
[  0] ROOT: No image
```

NOVA / Server: 3rd Gen Intel Xeon-SP (c6i.metal)

```
NOVA Microhypervisor #8b6c28a (x86_64): Feb  1 2024 10:56:51 [gcc 13.2.0]

[ -1] XSDT: 0x74bfc0e8 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:    188 (0xf5) ok
[ -1] FACP: 0x74bf7000 OEM:INTEL  TBL:INTEL ID REV: 6 LEN:    276 (0xb4) ok
[ -1] SSDT: 0x74bfb000 OEM:INTEL  TBL:xh_wccrb REV: 2 LEN:   2329 (0xc8) ok
[ -1] SSDT: 0x74bfa000 OEM:INTEL  TBL:RAS_ACPI REV: 2 LEN:   1864 (0x19) ok
[ -1] SSDT: 0x74bf9000 OEM:INTEL  TBL:ADDRXLAT REV: 2 LEN:   1861 (0x4f) ok
[ -1] ERST: 0x74bf8000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:    560 (0xdd) ok
[ -1] HMAT: 0x74bf6000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:    384 (0x4e) ok
[ -1] HPET: 0x74bf5000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:     56 (0x51) ok
[ -1] MCFG: 0x74bf4000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:     60 (0xde) ok
[ -1] MSCT: 0x74bf3000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:    100 (0xff) ok
[ -1] APIC: 0x74bbb000 OEM:INTEL  TBL:INTEL ID REV: 4 LEN:   1118 (0x23) ok
[ -1] SLIT: 0x74bba000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:     48 (0x44) ok
[ -1] SRAT: 0x74bb6000 OEM:INTEL  TBL:INTEL ID REV: 3 LEN:  12848 (0x2a) ok
[ -1] OEM4: 0x6e36d000 OEM:INTEL  TBL:CPU  CST REV: 2 LEN: 802145 (0x1b) ok
[ -1] OEM1: 0x74b2d000 OEM:INTEL  TBL:CPU EIST REV: 2 LEN: 558729 (0x46) ok
[ -1] OEM2: 0x74b09000 OEM:INTEL  TBL:CPU  HWP REV: 2 LEN: 143409 (0x6d) ok
[ -1] SSDT: 0x6e2b9000 OEM:INTEL  TBL:SSDT  PM REV: 2 LEN: 242286 (0x6c) ok
[ -1] HEST: 0x74b08000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:    380 (0xac) ok
[ -1] DMAR: 0x74b07000 OEM:INTEL  TBL:INTEL ID REV: 1 LEN:    368 (0x9c) ok
[ -1] SPCR: 0x74b06000 OEM:INTEL  TBL:INTEL ID REV: 2 LEN:     80 (0x31) ok
[ -1] ACPI: Version 6.2 Profile 4 Features 0x4a5 Boot 0x10
[ -1] FACS: Hardware 0x0 Flags 0x0 Wake 0x0/0x0
[ -1] SPCR: Console 8000:0000 (0:0xea600000:8:0)
[ -1] PCIE: 0x80000000 Segment 0x0000 Bus 0x00-0xff
[ -1] PCIE: 8086:09a2 08-80-00 D0 PCIE              0000:00:00.0
[ -1] PCIE: 8086:09a4 08-80-00 D0 PCIE              0000:00:00.1
[ -1] PCIE: 8086:09a3 08-80-00 D0 PCIE              0000:00:00.2
[ -1] PCIE: 8086:0998 06-00-00 D0 PCIE              0000:00:00.4
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.0
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.1
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.2
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.3
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.4
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.5
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.6
[ -1] PCIE: 8086:0b00 08-80-00 D0 PCIE PMI FLR      0000:00:01.7
[ -1] PCIE: 8086:09a6 08-80-00 D0                   0000:00:02.0
[ -1] PCIE: 8086:09a7 08-80-00 D0                   0000:00:02.1
[ -1] PCIE: 8086:3456 13-00-00 D0                   0000:00:02.4
[ -1] PCIE: 8086:a1ec ff-00-00 D0      PMI          0000:00:11.0
```

```
[ 48] CPST: 35-29-8 using 35 (HWP)
[ 48] RDTA: L3:15 L2:0 MB:15
[ 48] FPUC: State:0xe7 Size:2432
[ 48] VMCS: Revision:0x13 (0xfff9fffe:0x95d1def:0x0)
[ 48] CORE: 01:16.0 6:6a:6:0 [d0003b9] Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz
[ 90] APIC: LOC:0x35 VER:0x15 SUP:1 LVT:0x6 (x2APIC DL Mode)
[ 90] CCST: C6 C1 (0x1a0084005a:0x14008402)
[ 90] CPST: 35-29-8 using 35 (HWP)
[ 90] RDTA: L3:15 L2:0 MB:15
[ 90] FPUC: State:0xe7 Size:2432
[ 90] VMCS: Revision:0x13 (0xfff9fffe:0x95d1def:0x0)
[ 90] CORE: 00:26.1 6:6a:6:0 [d0003b9] Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz
[ 21] APIC: LOC:0x2a VER:0x15 SUP:1 LVT:0x6 (x2APIC DL Mode)
[ 21] CCST: C6 C1 (0x1a0084005a:0x14008402)
[ 21] CPST: 35-29-8 using 35 (HWP)
[ 21] RDTA: L3:15 L2:0 MB:15
[ 21] FPUC: State:0xe7 Size:2432
[ 21] VMCS: Revision:0x13 (0xfff9fffe:0x95d1def:0x0)
[ 21] CORE: 00:21.0 6:6a:6:0 [d0003b9] Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz
[122] APIC: LOC:0xb5 VER:0x15 SUP:1 LVT:0x6 (x2APIC DL Mode)
[122] CCST: C6 C1 (0x1a0084005a:0x14008402)
[122] CPST: 35-29-8 using 35 (HWP)
[122] RDTA: L3:15 L2:0 MB:15
[122] FPUC: State:0xe7 Size:2432
[122] VMCS: Revision:0x13 (0xfff9fffe:0x95d1def:0x0)
[122] CORE: 01:26.1 6:6a:6:0 [d0003b9] Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz
[127] APIC: LOC:0xbf VER:0x15 SUP:1 LVT:0x6 (x2APIC DL Mode)
[127] CCST: C6 C1 (0x1a0084005a:0x14008402)
[127] CPST: 35-29-8 using 35 (HWP)
[127] RDTA: L3:15 L2:0 MB:15
[127] FPUC: State:0xe7 Size:2432
[127] VMCS: Revision:0x13 (0xfff9fffe:0x95d1def:0x0)
[127] CORE: 01:31.1 6:6a:6:0 [d0003b9] Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz
[108] APIC: LOC:0x99 VER:0x15 SUP:1 LVT:0x6 (x2APIC DL Mode)
[108] CCST: C6 C1 (0x1a0084005a:0x14008402)
[108] CPST: 35-29-8 using 35 (HWP)
[108] RDTA: L3:15 L2:0 MB:15
[108] FPUC: State:0xe7 Size:2432
[108] VMCS: Revision:0x13 (0xfff9fffe:0x95d1def:0x0)
[108] CORE: 01:12.1 6:6a:6:0 [d0003b9] Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz
[  0] TIME: 29924ms 0ms/0ms
[  0] ROOT: No image
```

NOVA Microhypervisor: Innovation Timeline

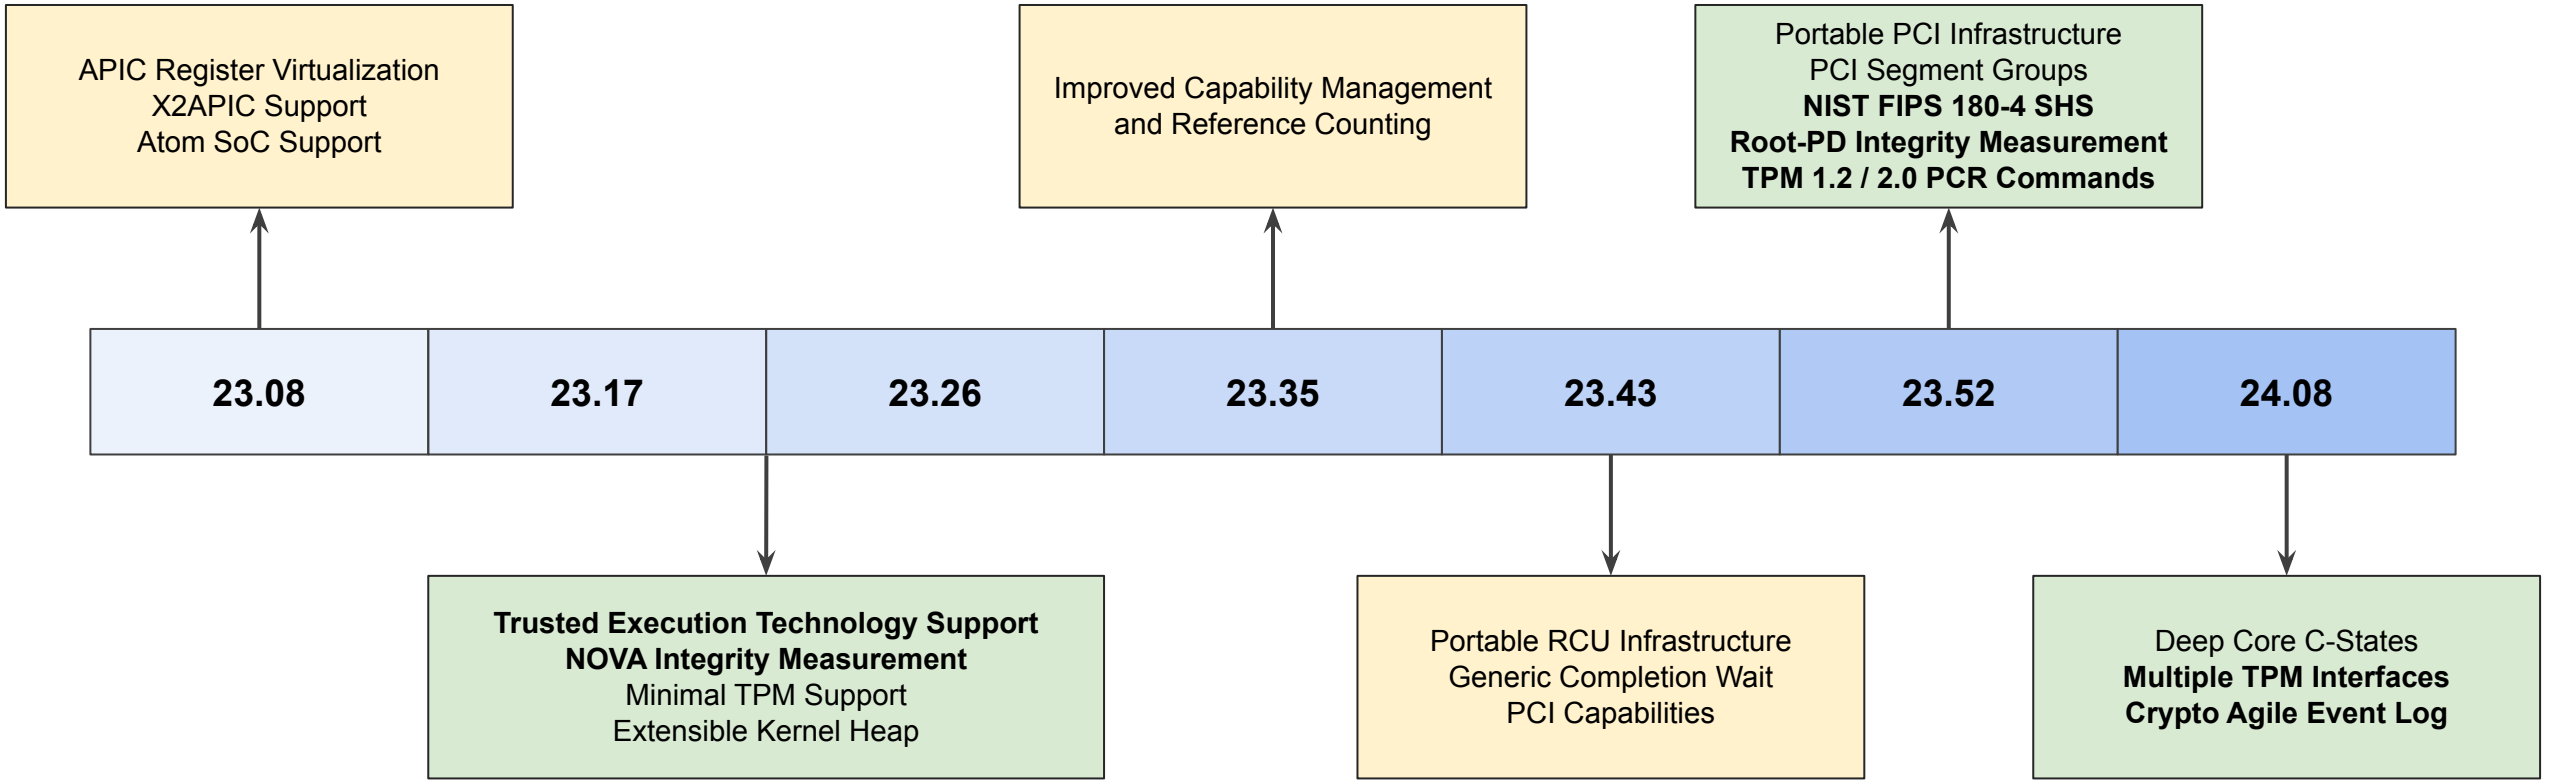

NOVA release cadence is ~2 months

What Problem Does Trusted Computing Solve?

- Once you have a formally verified software stack
  - and a compiler that produced a qualified set of binaries for the target architecture
- How do you ensure that some computer is running **those** binaries
  - and not some other (malicious) software instead
  - before you entrust that computer with your data or secrets
- In other words, how can you
  - either restrict the software that a computer will launch
  - or determine what software has been launched on a computer

Verified Boot: Static Root of Trust

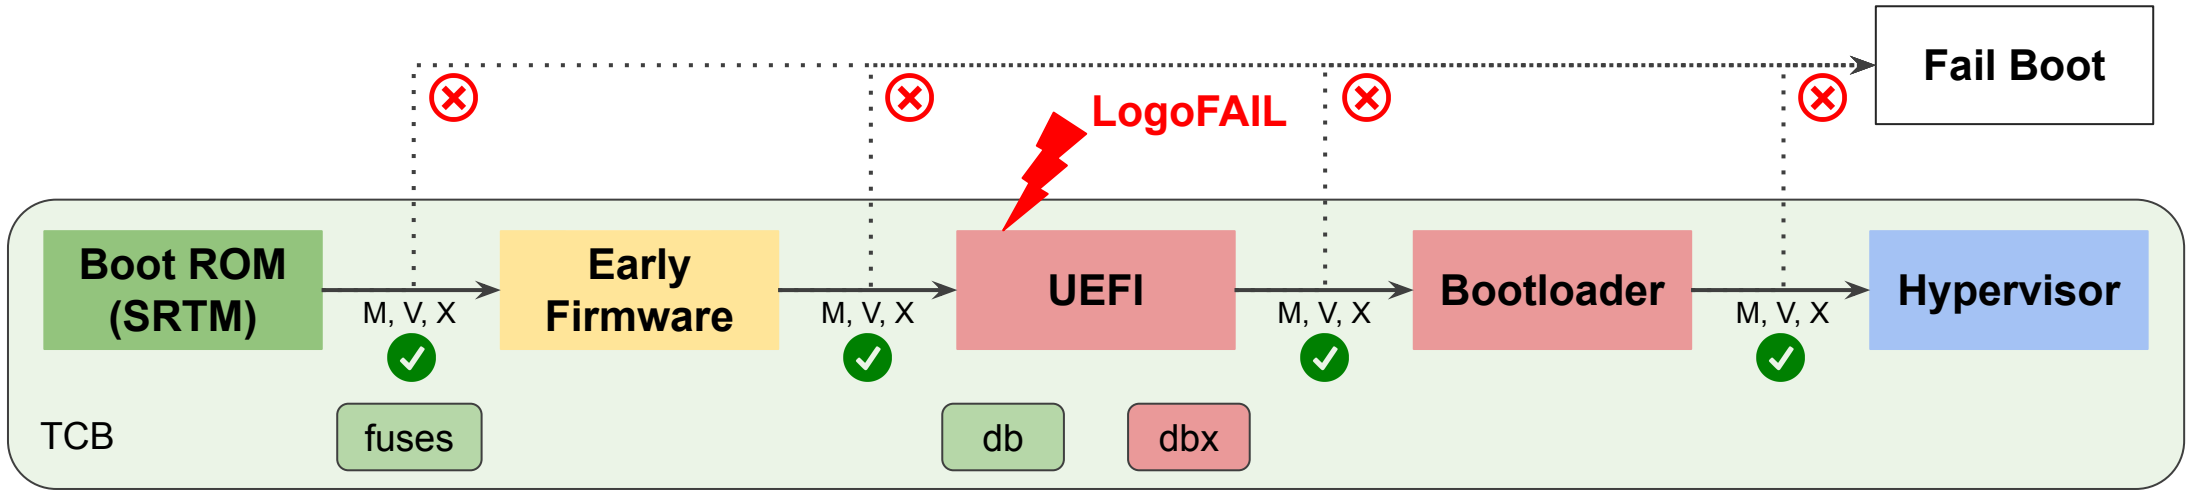

- Boot policies are enforced during the boot process
- Starting with the Core Root of Trust for Verification, the currently executing module verifies the integrity of the next module against a boot policy (e.g. UEFI db/dbx) ⇒ Chain of Trust
- Integrity measurement is a cryptographic hash ⇒ unique + indicative to changes in the module

Measured Boot: Static Root of Trust

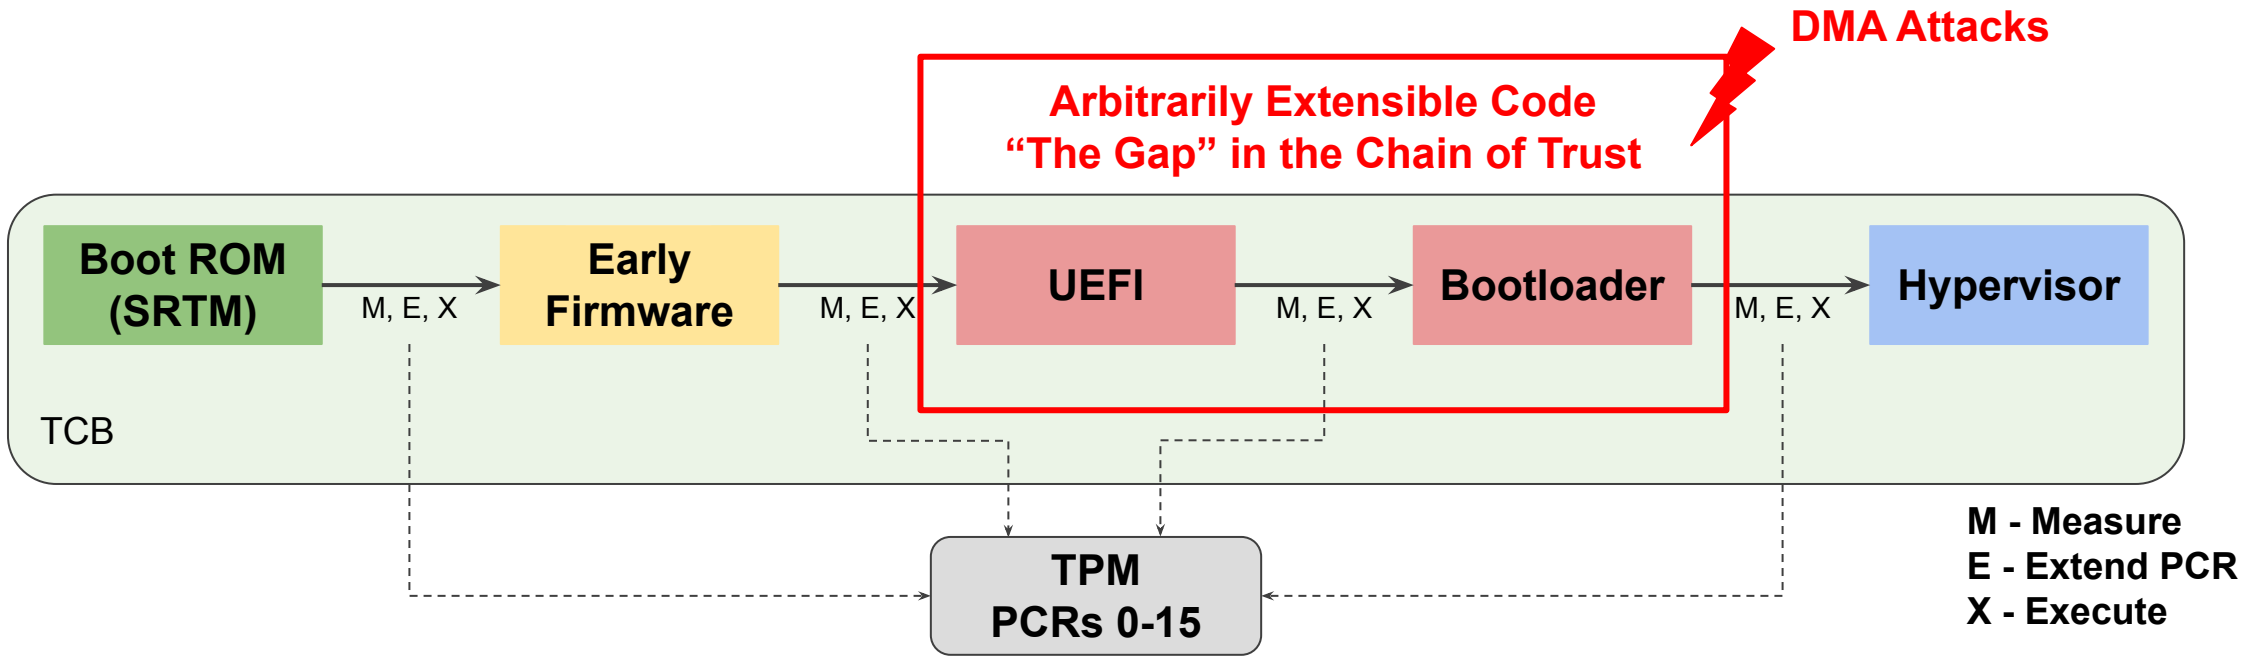

- Integrity measurements are extended into TPM PCRs during the boot process
- Starting with the Core Root of Trust for Measurement, the currently executing module extends the launch integrity measurement for the next module into the TPM

Measured Boot: Dynamic Root of Trust

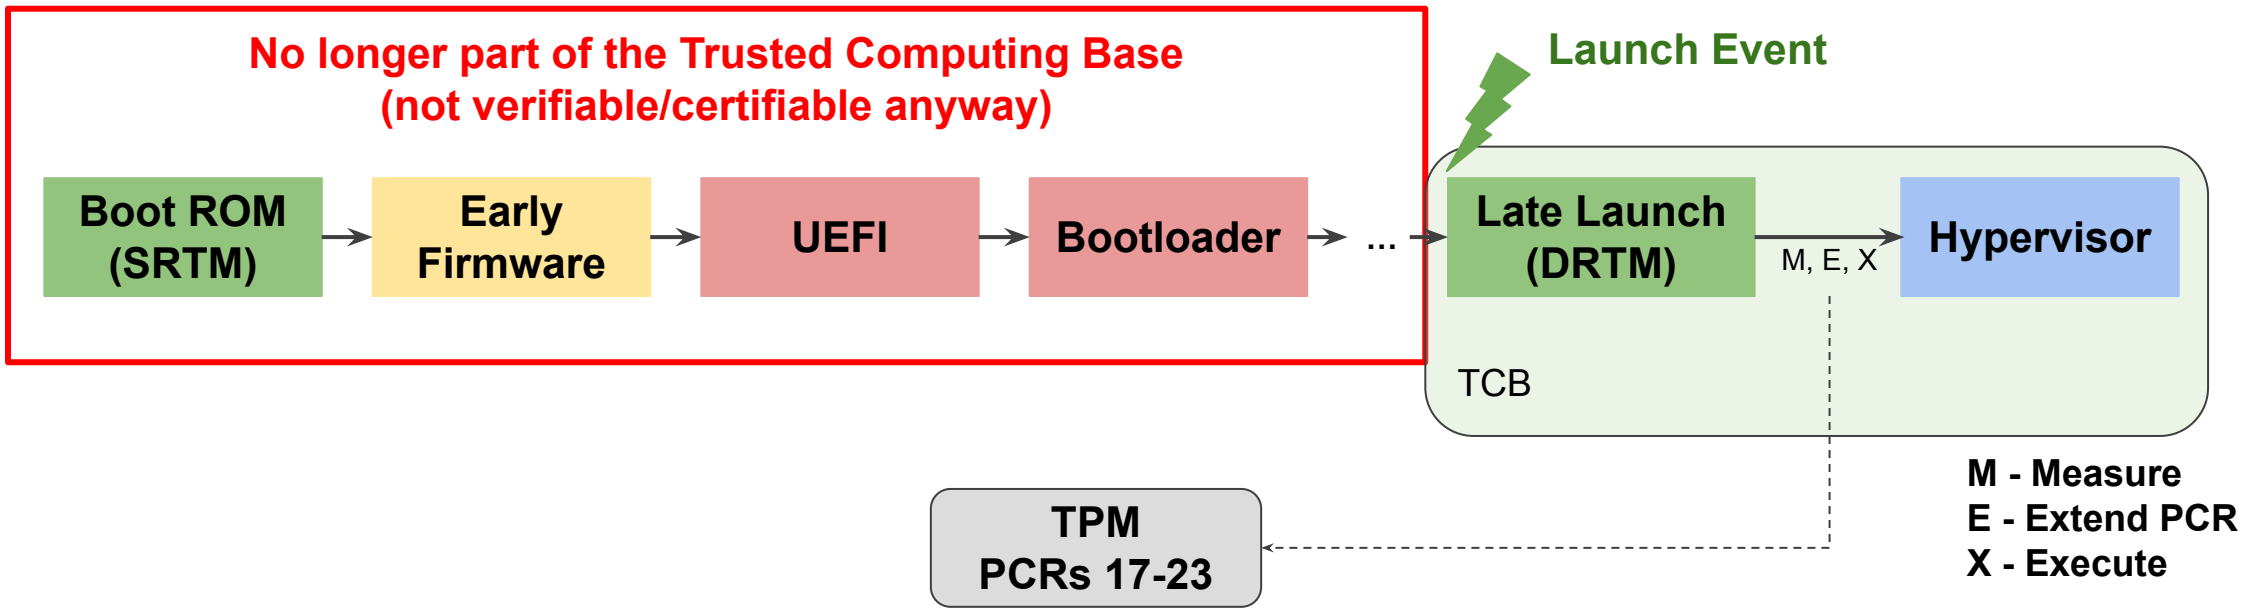

- DRTM Flow lets system boot into an untrustworthy state (initially)
  - Measured Launch later “resets” system into a trustworthy safe state
  - Takes control of all CPUs and forces them down a protected and measured code path

Intel Trusted Execution Technology (TXT / CBnT)

- Provides a Dynamic Root of Trust (DRTM)
- Prerequisites
  - CPU support (SMX features)
  - TXT-capable chipset (DMA protection)
  - TPM 2.0 (preferably) or 1.2
  - SINIT Authenticated Code Module (ACM)
- Use Cases
  - Remote Attestation (via TPM Quote)
  - Local Attestation (via Launch Control Policy)

GETSEC[SENTER] Late Launch Sequence

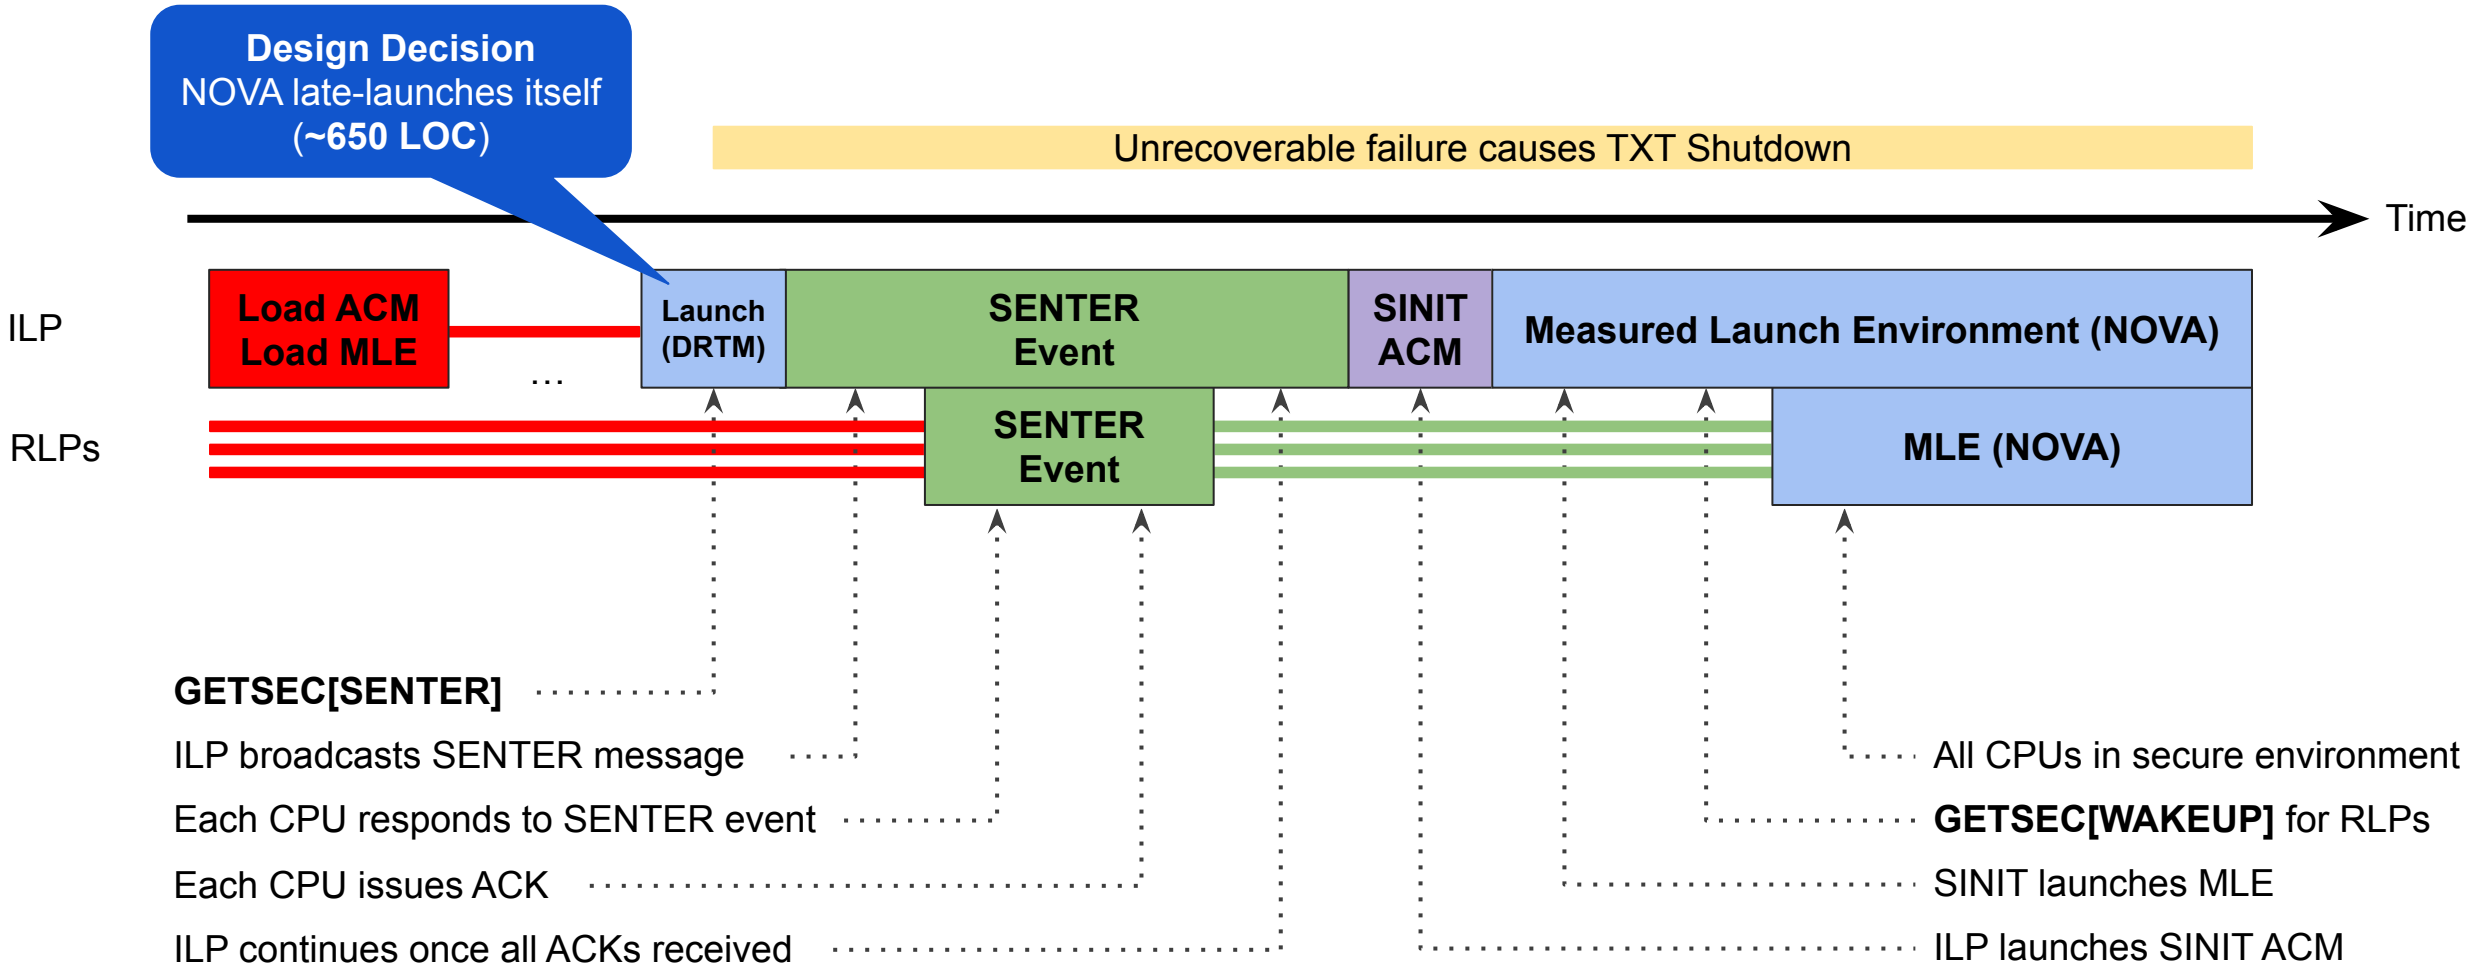

Trusted Platform Module (TPM)

Image Source:
Infineon Technologies

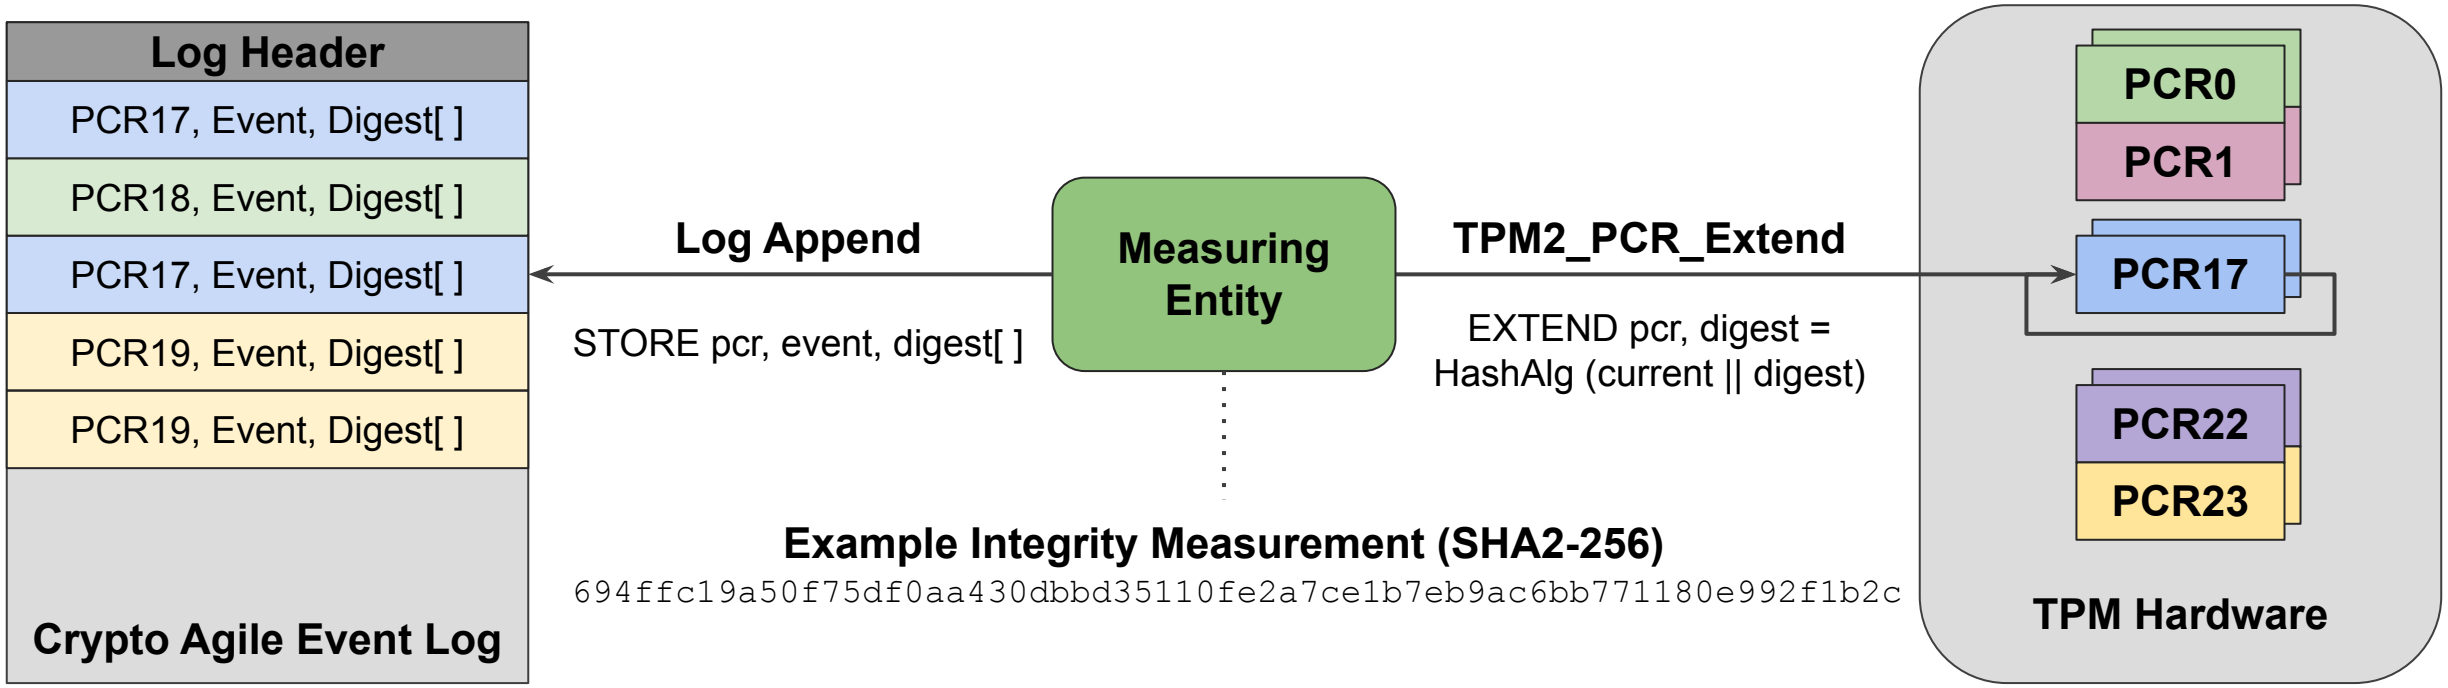

A verifier can use the crypto agile event log to recompute/validate the composite value in each PCR

Integrity Measurement (NOVA)

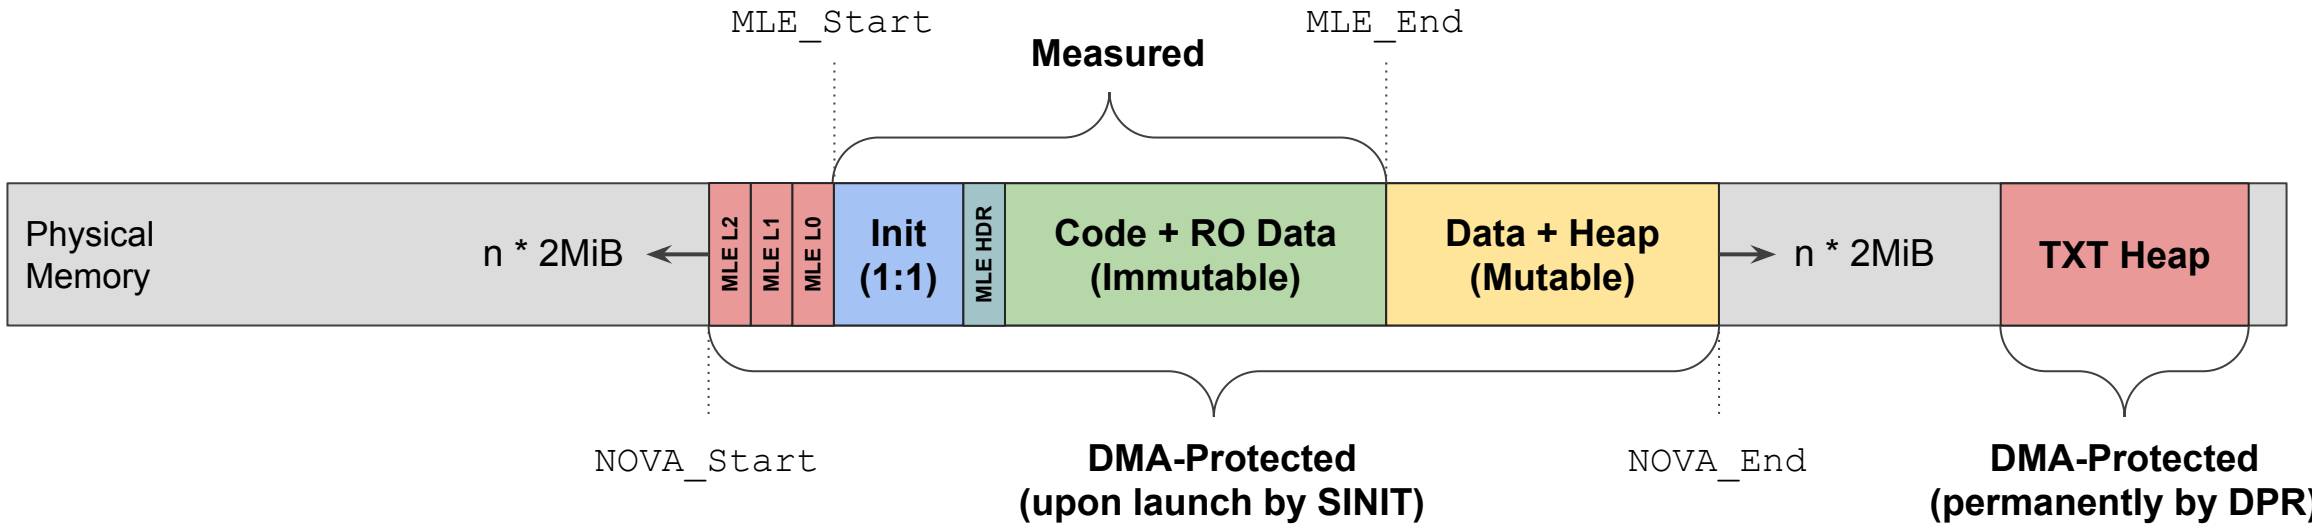

- What (not) to include in integrity measurement requires careful consideration
  - Sensitive to: modifications in NOVA measured region, command-line parameters
  - Insensitive to: hardware platform, memory map, multiboot modules, code self-patching
- Build system prints reference integrity measurements for NOVA

Extending the Chain of Trust into User-Mode

- Compute the Root-PD launch integrity measurement
  - Define the attestable region to measure
  - Compute the integrity measurement of that region
    - Either using TPM Hash/Event Sequence (maximum crypto agility)
    - Or using a SW implementation in NOVA (maximum performance)
- Drive the Trusted Platform Module
  - Different MMIO Interfaces
  - Different TPM Families
- Append integrity measurement to the TPM event log

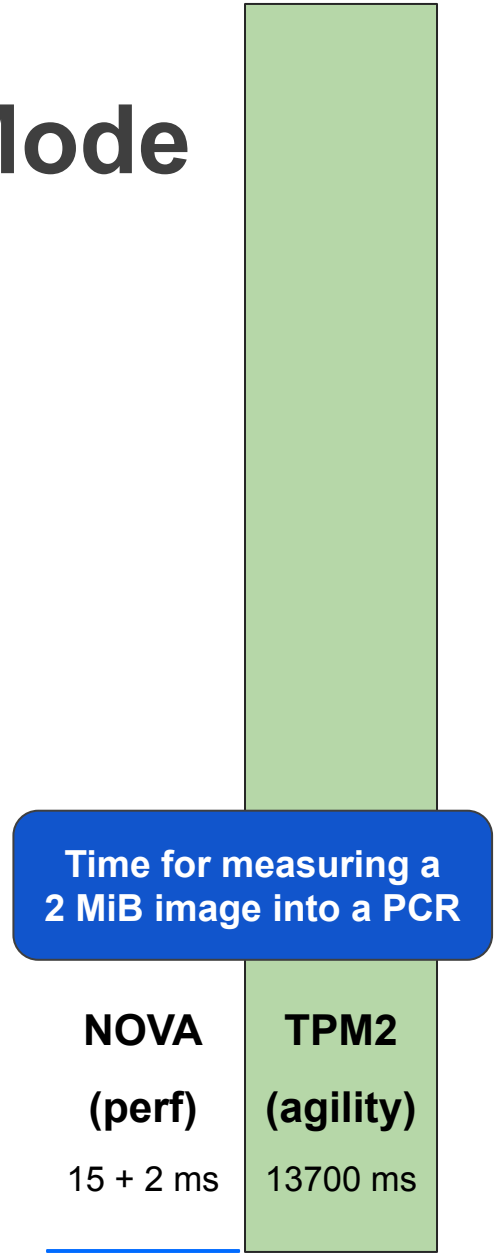

Integrity Measurement (Root Protection Domain)

- NOVA must measure Root-PD **before** launching it
  - Attestable region must be conveyed to NOVA without any invocations from Root-PD
  - Using ELF PHDRs (first non-writable PT_LOAD segment)
- Hash Computation
  - Digests computed by NOVA in C++ (for all supported hash algorithms)
  - Subsequently extended into PCR 19 (for all PCR banks)
- NOVA implements the FIPS 180-4 Secure Hash Standard (~130 LOC)
  - SHA1: -160
  - SHA2: -224, -256, -384, -512

Trusted Platform Module Infrastructure in NOVA

- Supports all TPM interface types (FIFO and CRB) (~250 LOC)
- Supports relevant command **subset** (Family 1.2 and 2.0) (~500 LOC)
  - Determine TPM capabilities (PCRs, Algorithms, …)
  - Perform PCR operations
- TPM localities control PCR access
  - Loc 2 belongs to NOVA Microhypervisor
  - Loc 1 belongs to User-Mode Environment
  - Loc 0 is for Legacy Use

Locality	Usage	Can Extend	Can Reset	Next Stage
4	**CRTM**	0-18, 23	17-22	⇒ **PCR 17**
3	**SINIT ACM**	0-20, 23	16, 20-23	⇒ **PCR 17**
2	**NOVA**	0-23	16, 20-23	⇒ **PCR 19**
1	**Root PD**	0-16, 20, 23	16, 23	⇒ **PCR 20**

Confidential & Trusted Computing Building Blocks

- **Availability**
  - Cache & Memory Bandwidth Allocation Technology (CAT/CDP/MBA) - since 22.26
- **Integrity**
  - Control-Flow Enforcement Technology (CET IBT+SSS) - since 22.17
- **Confidentiality**
  - Total Memory Encryption with Multiple Keys (TME-MK) - since 22.52
- **Measured Launch & Attestation**
  - Trusted Execution Technology (TXT/CBnT) - since 23.26

Questions and Discussion

The NOVA microhypervisor is licensed under GPLv2

Releases: https://github.com/udosteinberg/NOVA/tags

More Information: bedrocksystems.com and hypervisor.org