DC-XPI

DataCenter-ready eXtended Peripheral Interface

An Interconnect Proposal for Modular I/O (DC-MIO)

Mike Branch, Siamak Tavallaei - Google
Nilesh Dattani, Priya Raghu - Microsoft

2019: OCP Summit

AIC Attachment

IO Slot to CPU Board Cable Harness

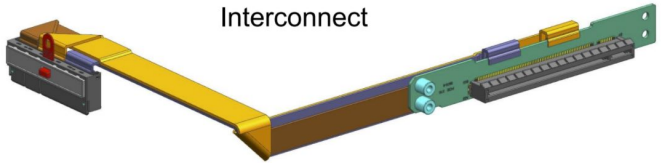

Open. Together.

For a successful Modular Building Block Architecture, we need:

- Compute Modules (CPU/Memory/IO) (**CMIOM**)
- IO & Accelerator Add-in Card Modules (**AIC**)
- Security, Control, and Management (**SCM**)
- Data-plane Control
- An Interconnect

OCP SUMMIT

Open. Together.

2021: This has evolved to:

DataCenter-ready Modular Hardware System (DC-MHS)

Datacenter-ready Modular Hardware System

An overview from: [OCP Server Project Monthly Call Presentation on DC-Stack](#) (5/26/2021)
for Enterprise, Hyperscale, and Edge datacenter

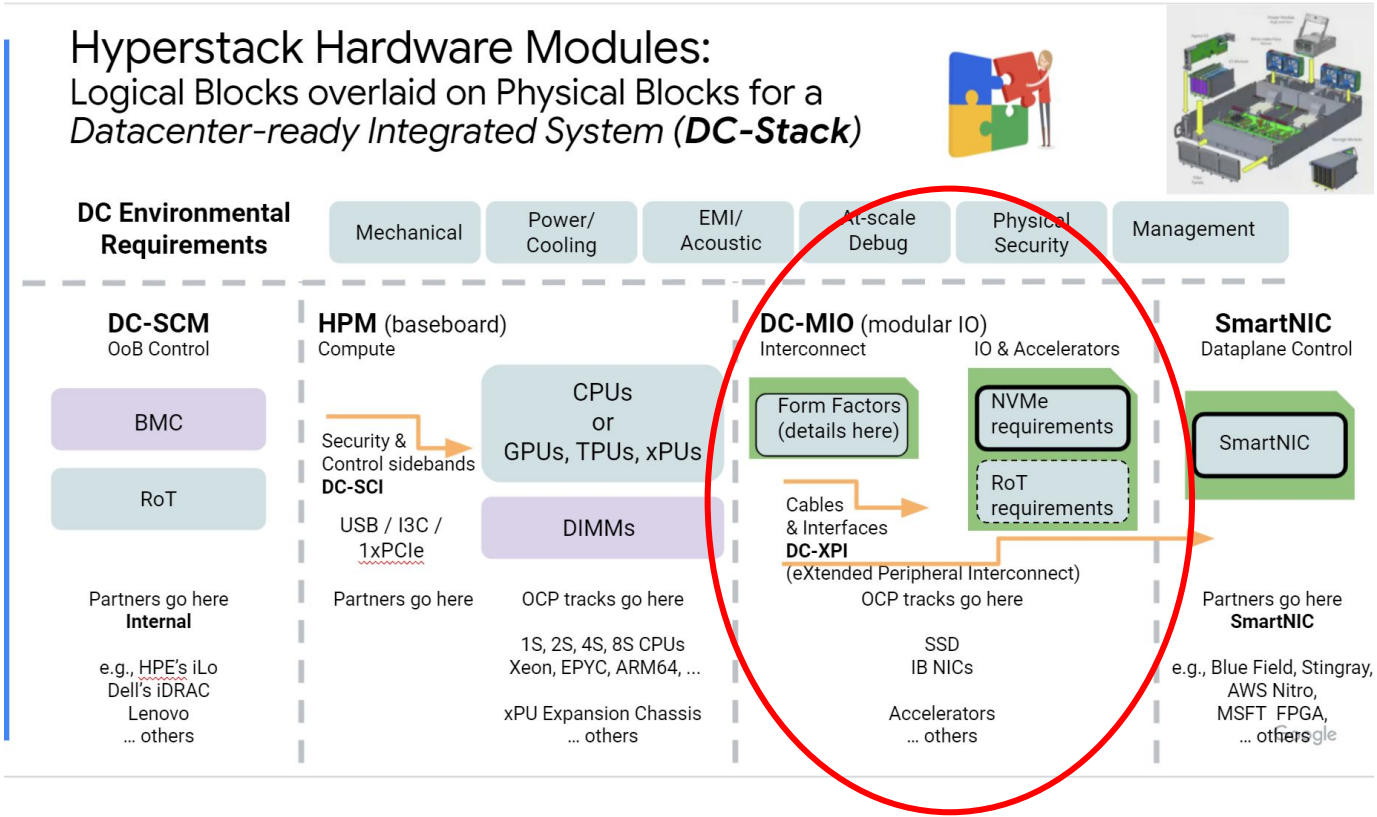

Why I/O Modularity?

At the high level:

- Interface speeds have been increasing
    - Increasing mobo material costs and/or
    - Increasing need for re-timers
- Higher power peripherals (requiring additional cabling)
- Increasing # of peripheral shapes to support (CEM, U.2, EDSFF, custom, ...)
- Desire for “pay-as-you-go” addition of peripherals
- Increasing # of server platforms to validate

DataCenter-ready Modular I/O (DC-MIO)

- Packaging approach that separates the motherboard (HPM[1]) from the I/O peripherals
- Allows high-speed I/O connector(s) near the CPU(s)
- Uses I/O Adapters to connect peripherals to the HPM
- Reduces motherboard size & cost
- Allows for cabled and riser-style I/O Adapters
    - Cabled I/O adapters may eliminate need for retimers
- Accommodates multiple peripheral form factors
- I/O Adapters can be installed as-needed based on tray config

[1] Host Processor/Memory Module

Implementation Goals

How should this modular interface be implemented?

Goals:

- A high-speed, high-density connector
- A high-volume connector with multiple sources
- Cable and riser-card support
- Support for x16 (not too concerned with optimizing for smaller width connectors)
- Support higher-power (12V) peripherals without additional cables
- Support a robust set of sideband interfaces
- Re-use existing high-volume connector and pinout if possible
- Support flexible mounting orientations: vertical/horizontal/coplanar (1U/2U/...)

An Implementation

DataCenter-ready eXtended Peripheral Interface (**DC-XPI**)

- SFF-TA-1002 4C+ connector provided the desired speed, density and pin count
    - PCIe Gen6, 0.6mm/<3” length, x16 + sidebands
- Connector already has volumes being driven by OCP NIC & DC-SCM
- Allows for cabled and riser-style I/O adapters
- Created a pinout that supports high power (150W) peripheral(s)
    - Supports 2x 75W CEM cards
- Optional (separate) auxiliary power block to support up to 400W peripheral(s)
- Rich set of sideband interfaces including USB2, USB3, UART, I2C
- Supports individual Presence Detect for I/O Adapter and Peripheral

A New Pinout?

Several existing pinout/connector options:

- EDSFF / PECFF (4C)
- PECFF-HP-12V (4C)
- OCP NIC 3.0 / PECFF (4C+)

But none that supports:

- High power (150W) peripherals without additional power cables
  -and-
- A rich set of sideband interfaces including USB2, USB3, and UART

Implementation

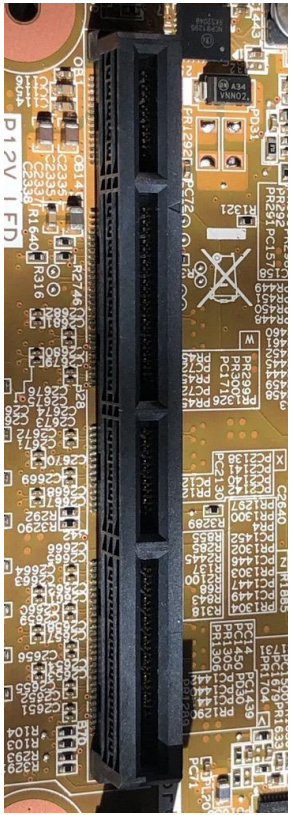

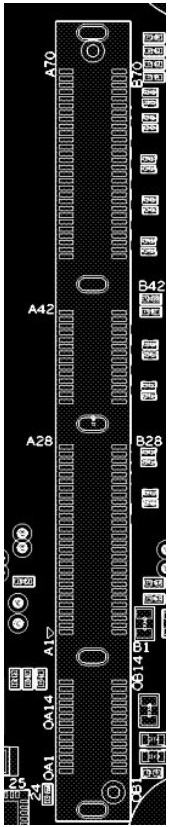

Implementation (cont'd)

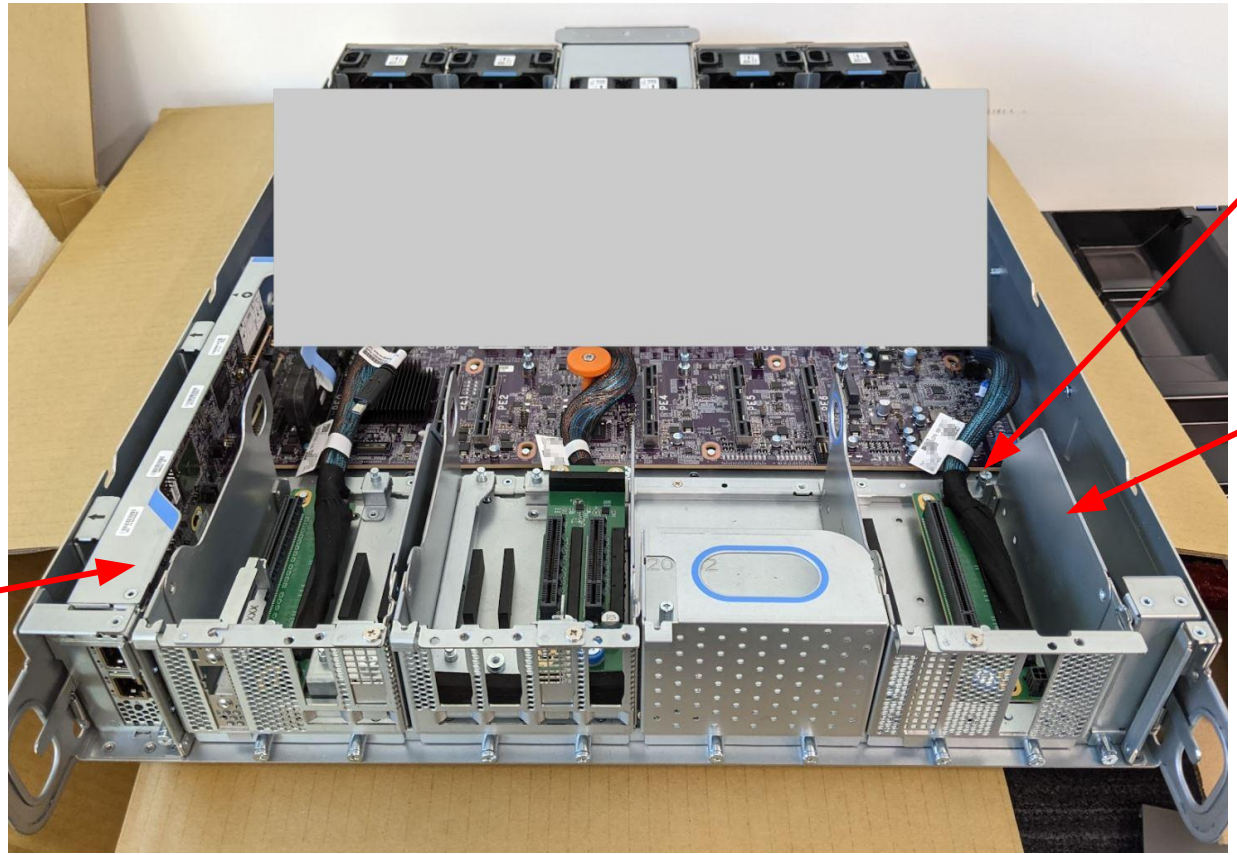

Example of a front I/O server using Modular I/O w/ vertical DC-XPI connectors (and DC-SCM).

Implementation (cont'd)

1x16 CEM cabled I/O Adapter

2x8 CEM cabled I/O Adapter

Allows for riser-based I/O Adapters, as well

Multiple vertical DC-XPI connectors across front of HPM

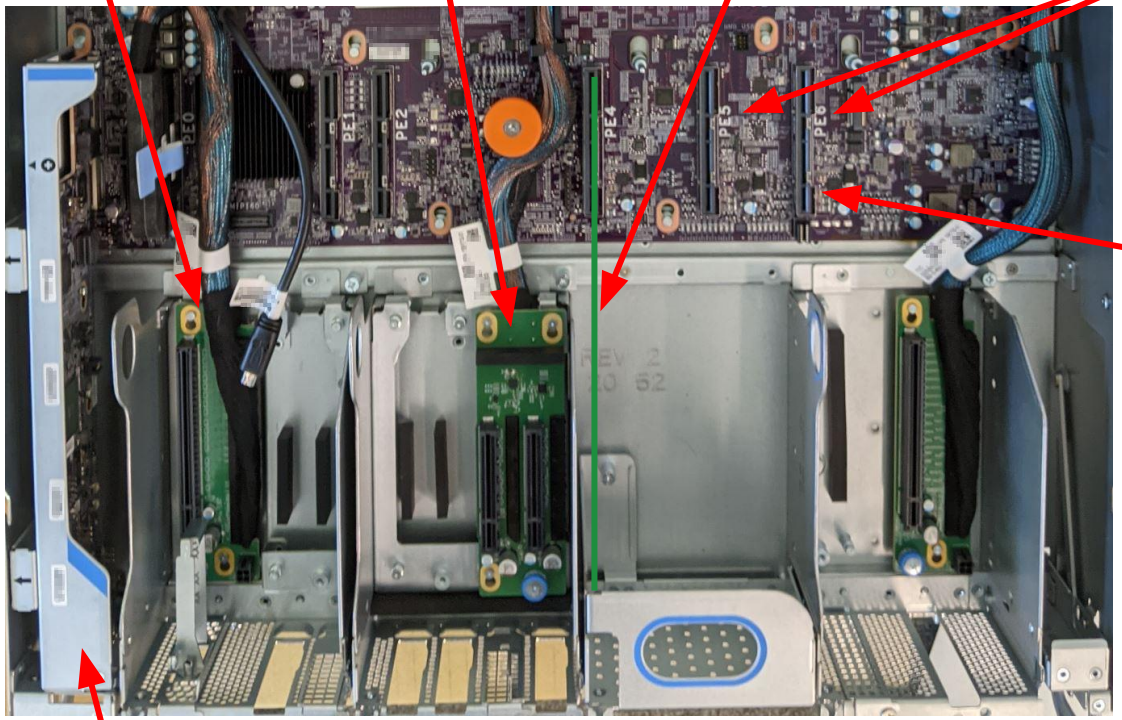

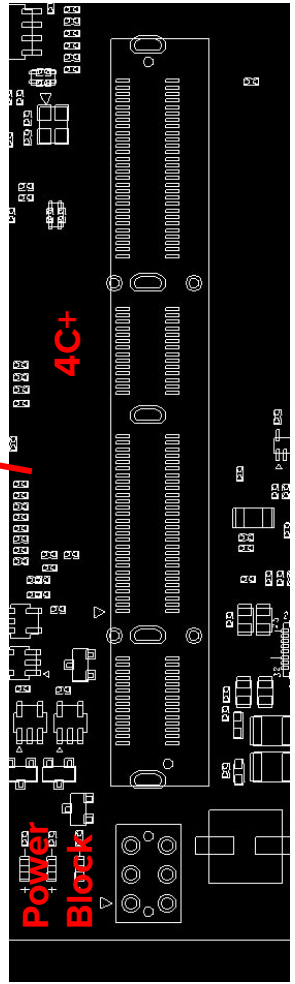

DC-SCM

(top view)

Implementation (cont'd)

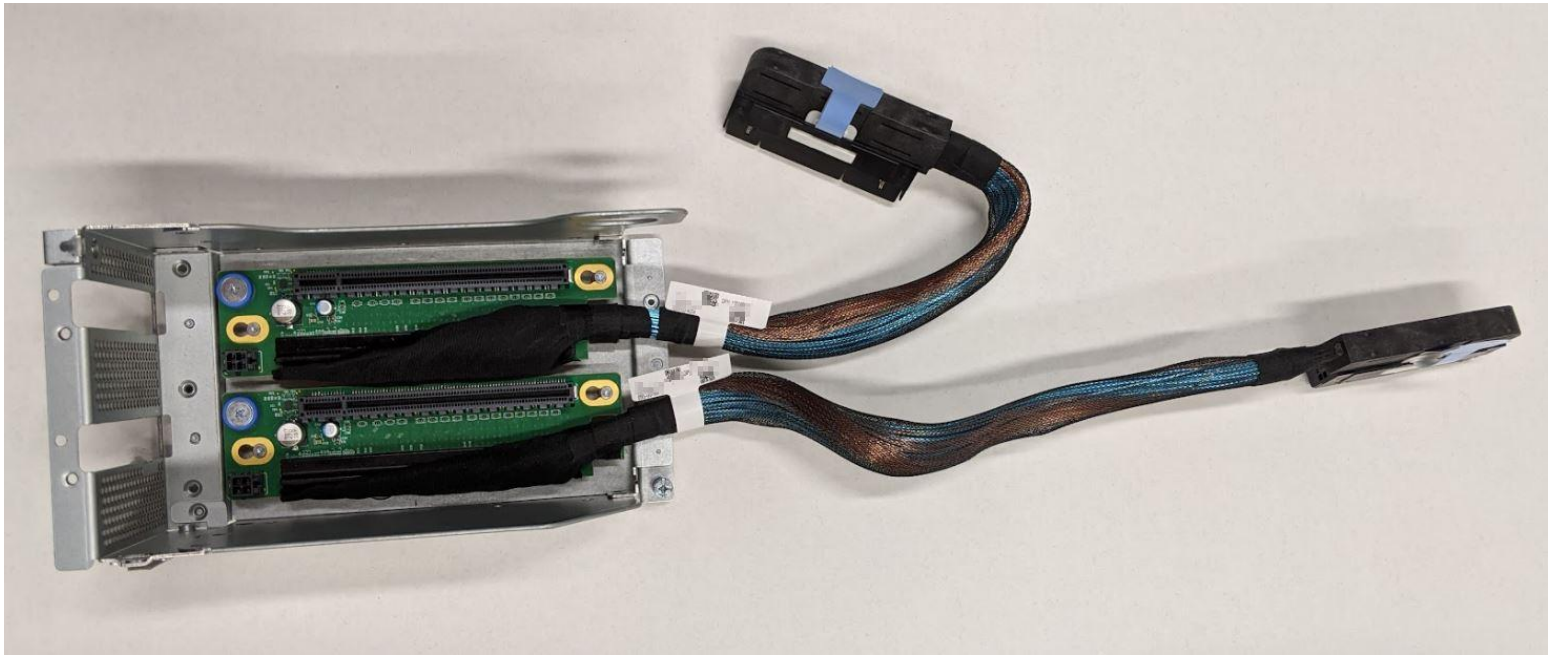

Two 1x16 Cabled CEM I/O Adapters in an I/O Module
(top view)

DC-XPI Spec Status / Next Steps

The DC-XPI 1.0 spec has been largely completed for productization in 2022.

Similar to DC-SCM 1.0, we hope to gather support and feedback from OCP members which could lead to a second iteration of the spec, i.e., DC-XPI 2.0.

We are targeting the **DC-XPI** 2.0 spec for use in 2023+ servers, coincident with the **DC-SCM** 2.0 and **DC-MHS** 1.0 specs for **DC-Stack** 1.0

We plan to present DC-XPI at the OCP Summit Workshops along with DC-Stack!

Backup

Alignment with other OCP Activities

The following figure depicts where Datacenter-ready Integrated System (DC-Stack) falls within the continuum from DC-SCM through the datacenter-level Cloud Service Model initiative within OCP.

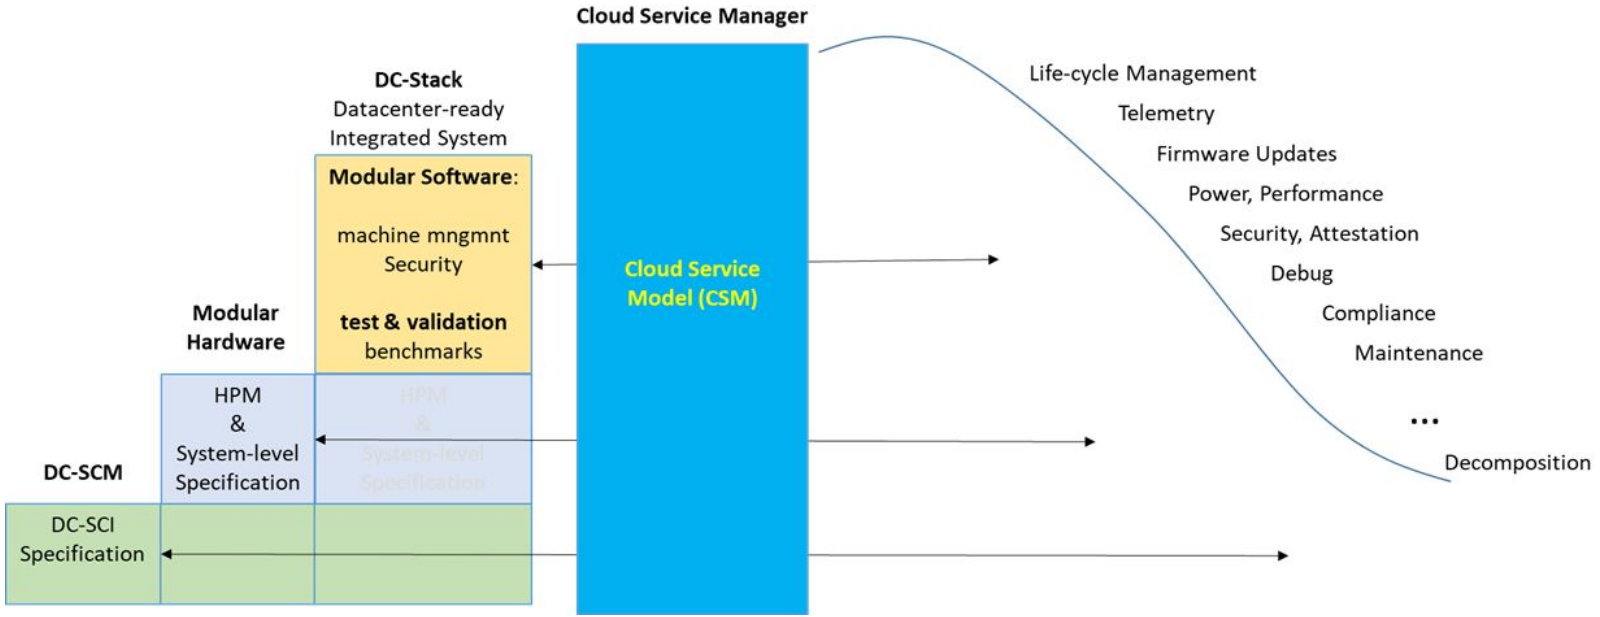

Datacenter-ready Integrated System *(DC-Stack)*

A convolution of many essential ingredients

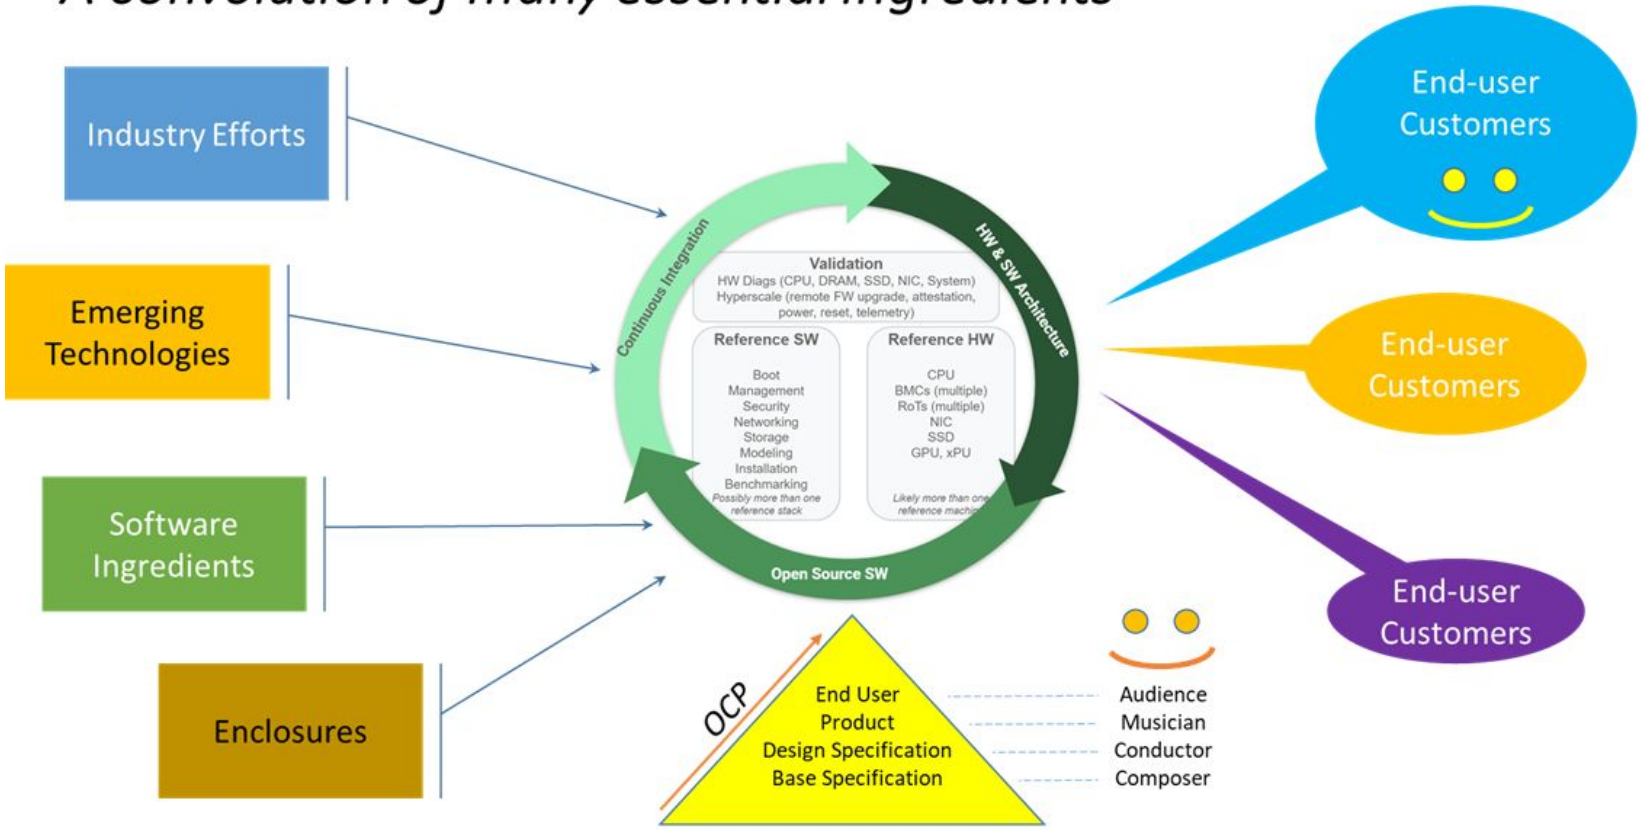

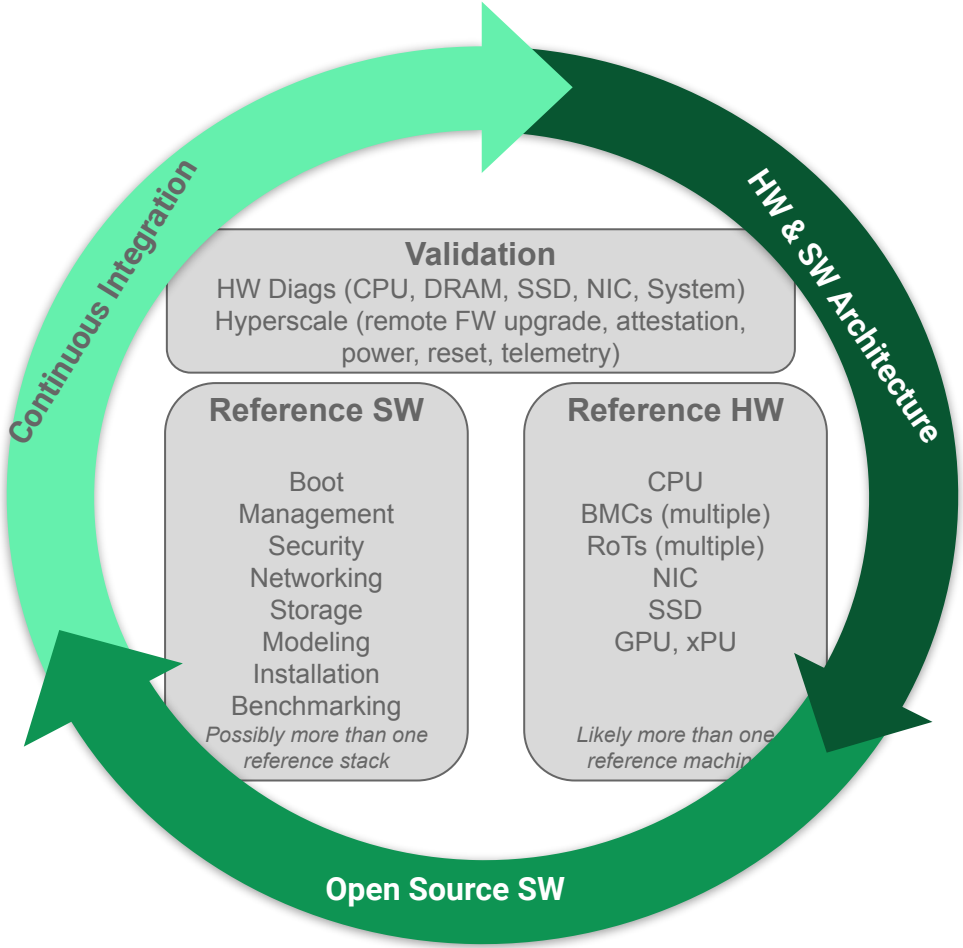
Continuous Integration
HW & SW Architecture
Open Source SW
Validation
HW Diags (CPU, DRAM, SSD, NIC, System)
Hyperscale (remote FW upgrade, attestation, power, reset, telemetry)
Reference SW
Boot
Management
Security
Networking
Storage
Modeling
Installation
Benchmarking
Possibly more than one reference stack
Reference HW
CPU
BMCs (multiple)
RoTs (multiple)
NIC
SSD
GPU, xPU
Likely more than one reference machin