

# OCP-T Spec Open Pod Update

Defining the next generation of OCP  
and RSD Open Pods

Jeff Sharpe

January, 2017



**Building Forward Together**



# Scope for OCP-T



Pod for Frame/appliance layouts

## **Current OCP-T Spec:**

- OCP-T frame, Power, Interconnect & Pod dimension spec submitted by Radisys:
  - Current spec defines the dimensions of 2 "Pods", ½ and Full Width 2U Pods
  - Pre-defines interfaces and DC Power requirements into Frame
  - Does not specify characteristics of the Pod internal component layout

## **ADLINK Proposed OCP-T Pod Spec:**

- Specs are used as a guideline for appliance delivery that is based on Modular Industrial Compute Architecture - CPU, Appliance, Future Switching (MR-IOV Plex and RRC as examples)
- Pod spec defines zones, what the zones are for and parameters for zones and use MICA as reference). Power, Mezzanine and interconnections between components are defined
- Both ½ width and Full Width POD Spec planned

# OCP-T Server Spec Overview



## Benefits

- Enables a common architecture for OCP-T suppliers to build a standard Central office product
- Provides both ½ Width and Full Width Pods for a multitude of options
- Defines the zones of each pod for operators to feel confident that the POD is open, however provides many options for POD use cases
- Enables the front part of the sled (Zone 5) and the Mezzanine (Zone 4) for Server, ARM, HW acceleration, Front Panel options, Storage, etc.
- Future Mezzanine enables ability to add multi-host controllers, MR-IOV capabilities, switching and HW acceleration for additional capabilities.
- Utilizes Radisys' best in class spec for Frame and Sled interconnection, power and physical dimensions.

# OCP Telecom Standardization based on MICA



## 19" OCP-T Frames



## "Pods" for OCP-T & OpenRack





1/2 Width



Full Width

## Reusable Sleds



## OCP-T Specification



### POD Spec:

- POD Zones & Layout
- POD Power
- POD I/O
- POD Environmental




# Open Pod Zone Specification




## POD Spec:

- Defines POD Zones, Dimensions & Layout for ½ and full width sleds
- POD Power requirements
- POD I/O options, Mezzanine and NIC
- POD Environmental for cooling and heat dissipation
- Provides interface connections between each zone
- Mezzanine interface for additional multi-hosted options
- Enables multiple Server options



# Proposed OCP-T “OpenPod” Layout

## OCP-T Pod Zones




## Definition for ½ width and full width OpenPods

- Fits into the OCP-T CG OpenRack spec utilizing MICA: PCIe Midplane, CPU and NIM sizes
- Dimensions and interconnects defined by OCP-T CG OpenRack Spec
- Defines interconnections between each zone above
- Utilizes OCP-T Power and Frame interconnect as defined by OCP-T CG OpenRack Spec
- Zones 1 – 3 Defines the foundation of the POD for system connections, Cooling, Power and Mid-Plane
- Zones 4 – 5 Defines the compute & optional interfaces connecting into the Pod base

# Full Width (2U) OCP-T Tray with 4 CPU sleds



Full Mezzanine enabling passive connections for POD switching. Multi-host controller, PCIe Switching within the POD as a single appliance  
Could be a complete storage sled as well.



1x OCP Power Bus Connector

8x Optical Data Port Connector


1x OCP Power Bus Connector

8x Optical Data Port Connector

# M.I.C.A. Plan of Intent (2017-2020)




## Sleds




Purley / FPGA / TickTock



5G Acceleration



NFV/SDN/IoT Acceleration



Continued increase of switching throughput



SSD/HDD Storage



GPGPU


## Systems



OCP Telecom



Hyper Converged Systems



Backplane Enhancements



Continued Platform s/w expansion



# Network Alliance program



Industrial leaders as part of our Modular Industrial Compute Alliance

## Alliance

Chassis, Sleds,  
Custom Systems

Application  
Partners

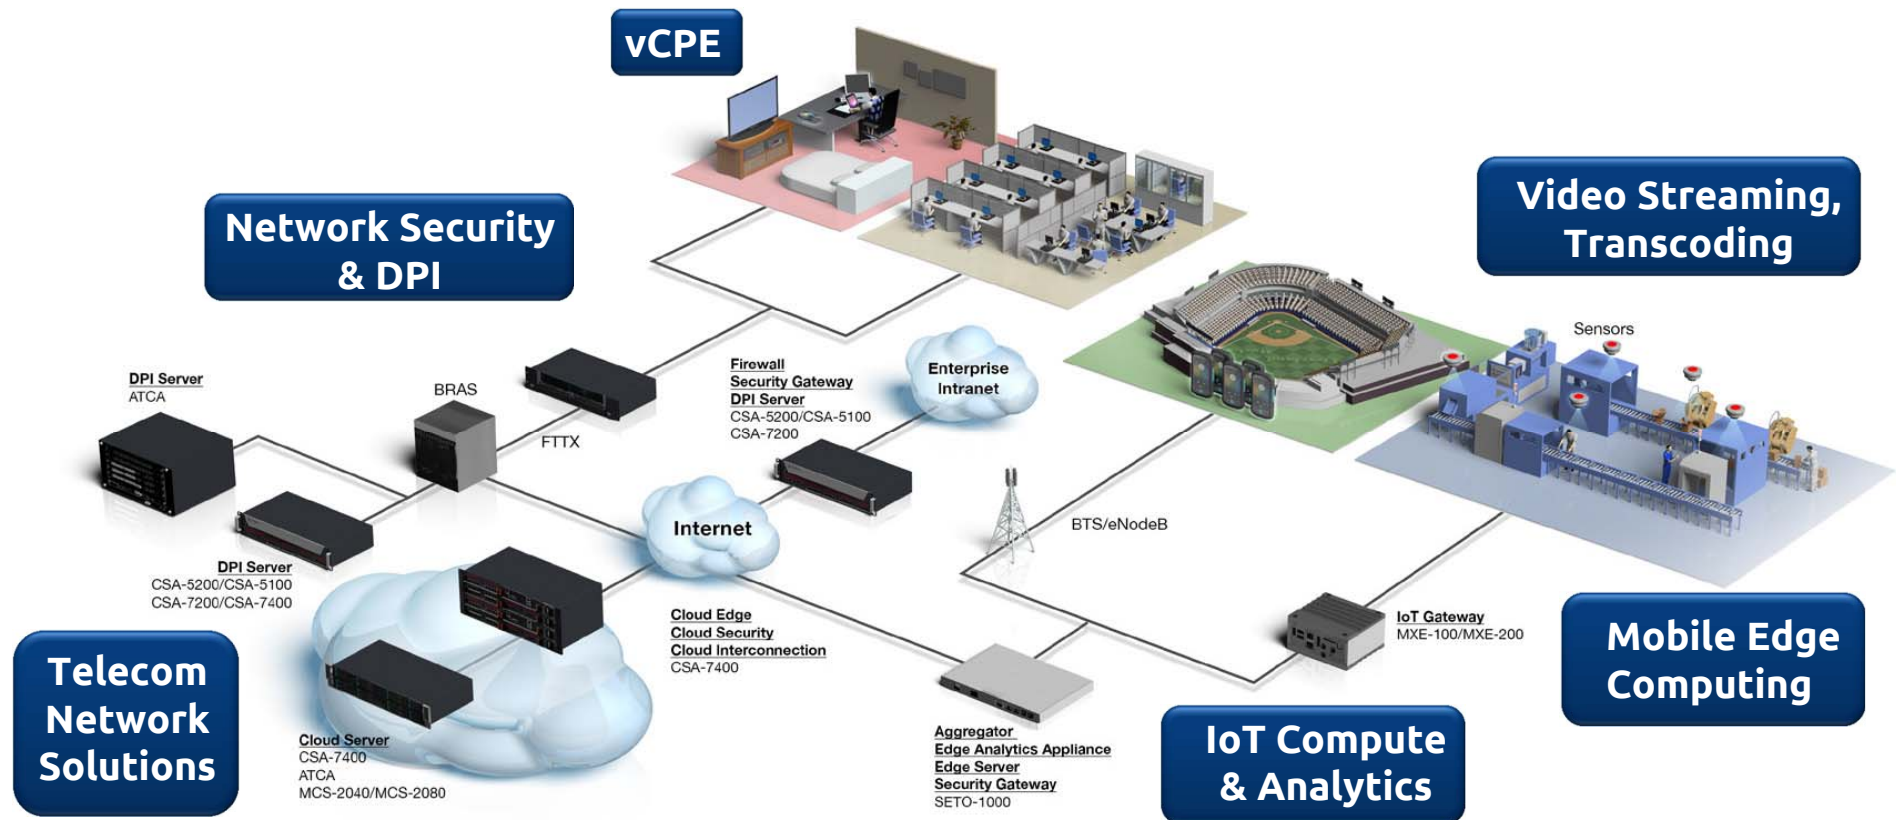
Integration, Testing  
& Support



Middleware


# Backup

# End-to-End Solutions Application Ready Platforms



# Mix & Match

## Rich system configuration choice




Mother Board  
MCN-2600

1x MCN-2600  
8x NIMs

4x MCN-2600  
2x MXN-3610/4100


24\*3.5' HDD

2x MCN-2600




CSA-7200  
8 NIMs,  
128 Eth Ports

Rich I/O




CSA-7400  
4 Compute Nodes,  
360G Throughput

High Compute  
Performance  
&Throughput



CSA-7x00  
24\*HDD+  
1\* CSA-7400

Large Storage, Storage,  
Compute, Switch Hyper  
Converged



OCP-T Compute Sled  
2 Compute Nodes,

# The Next Gen.- Network Communication Platform

With Flexible Modular Building Blocks

## Compute Nodes


- Broadwell, Skylake, Purley (Skylake + FPGA)
- ½ and ¼ width sleds for needed core density
- Flexibility: mix and match E3 and E5
- Integrated NFVi Software and Platform S/W




MCN-2600  
Broadwell



MCN-2610T  
Purley




MCN-1500  
(Greenlow GT4e)




MCN-1510  
Purley

## Switching Nodes


- Intel Red Rock Canyon & Broadcom PEX
- Extension trays for flexible I/O configs
- System Management and wide range of SW



MXN-3610  
(RRC)



MXN-3620  
(PCIe switch)



Planning: HW  
Accelerators &  
ARM

## I/O Nodes

- Optional Network Interface Modules
- Optical, copper, w and w/o bypass
- ¼ width size to support large I/O options



CSA-Z5C4F



CSA-ZBX10



CSA-Z5C2F



CSA-Z5C8F+



CSA-Z4X01

## Complete Systems


- 19" - 2U 4U and OCP-Telecom versions
- Data Center 2U
- AC/DC power, NEBS ready, customizable
- Multiple options based on required solution



CSA-  
7200/7400



MCS-2080



OCP



Custom



# Compute Node

## MCN-2600, MCN-1500

### MCN-2600

- 1/2 width
- Intel Grantley Refresh Platform
- 2x E5-2600 v3 CPU (1 CPU optional)
- 12 DIMM, 192G 1600MHz DDR4
- 2x 2.5" hot-swappable HDD
- PCIe G3 x 16 base connection, up to x64 expansion
- Redundant M.2 SSD
- Module management, IPMI I2.0
- Planned for 2017: MCN-2610T is Purley



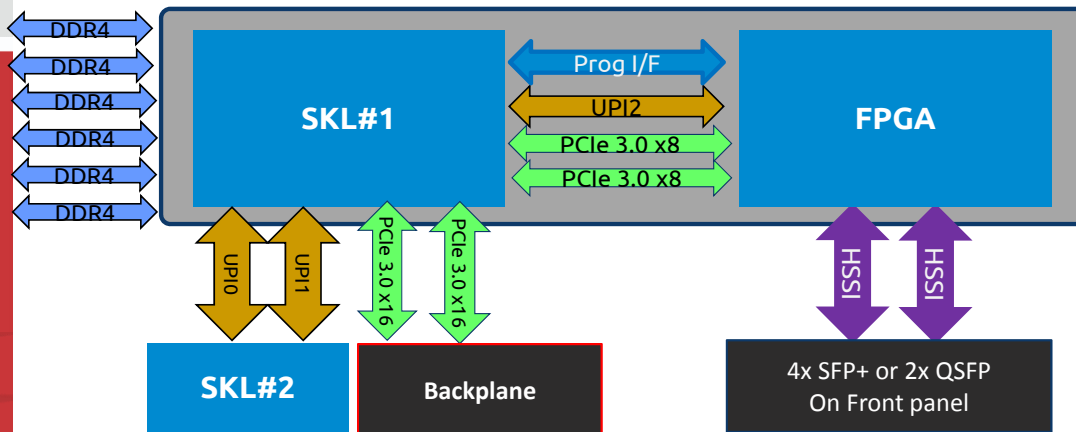
### MCN-1500

- 1/4 width
- Intel Skylake platform
- E3-1500 v5 CPU (GT4e Greenlow)
- 2 DIMM, 64 1600MHz DDR4
- 1x MiniSATA
- PCIe G3 x 8 backplane bandwidth
- Module management, IPMI 2.0



**Available Now**

## MCN-2610T: Xeon + FPGA + NVMe



- **Co-layout SKL non-FPGA & FPGA SKUs**
- **1.5X memory bandwidth with 6 vs. 4 memory channels**
- **Richer set of available IOs including more PCIe lanes, eSPI etc.**
- **Integrated 4x10GbE: Cost, power and area advantages vs. discrete Ethernet Si**
- **Single standard server development: Accelerate NFV transition**
- **Reduced total platform investment with application, control, and data plane workload consolidation**

<b>Cores</b>	Up to 28C with Intel® HT Technology	
<b>FPGA</b>	Altera® Arria 10 GX 1150	
<b>Socket TDP</b>	Shared socket TDP Up to 165W SKL & Up to <b>90W</b> FPGA	
<b>Socket</b>	Socket P	
<b>Scalability</b>	Up to 2S – with SKL-SP or SKL + FPGA SKUs	
<b>PCH</b>	Lewisburg: DMI3 – 4 lanes; 14xUSB2 ports Up to: 10xUSB3; 14xSATA3, 20xPCIe*3 New: Innovation Engine, 4x10GbE ports, Intel® QuickAssist Technology	
	<b>For CPU</b>	<b>For FPGA</b>
<b>Memory</b>	6 channels DDR4 RDIMM, LRDIMM, Apache Pass DIMMs	Low latency access to system memory via UPI & PCIe interconnect
	2666 1DPC, 2133, 2400 2DPC	
<b>Intel® UPI</b>	2 channels (10.4, 9.6 GT/s)	1 channel (9.6 GT/s)
<b>PCIe*</b>	PCIe* 3.0 (8.0, 5.0, 2.5 GT/s)	PCIe* 3.0 (8.0, 5.0, 2.5 GT/s)
	32 lanes per CPU Bifurcation support: x16, x8, x4	16 lanes per FPGA Bifurcation support: x8
<b>High Speed Serial Interface</b> (Different board design based on HSSI config)	N/A	2xPCIe 3.0 x8
		Direct Ethernet (4x10 GbE, 2x40 GbE, 10x10 GbE, 2x25 GbE)

# Compute Node

## MCN-2610T Purley & Purley w/ FPGA



### Highlight

- 1/2 width Skylake server CPU Node
- Intel Purley Platform
- Dual Skylake server CPU or dual Skylake server CPU + FPGA
- IPMI V2.0 Compliant

- Fully compliant to Modular Industrial Cloud Architecture CPU Module definition
- Same size as MCN-2600T, 1/2 width, 204x 430x 44mm (W x D x H)
- Co-layout Dual Skylake server CPU or dual Skylake server CPU + FPGA
- 2x 40G or 4x 10G Front IO direct to FPGA
- Lewisburg PCH
- 12x DDR4 RDIMMs, up to 384G
- Hot-swappable CPU Node
- 2x M.2 redundant SSD on board
- Front panel IO
  - . 2x RJ45 for management
  - . 1x RJ45 console port
  - . 1x VGA
  - . 4x SFP+ or 2x QSFP direct to FPGA only
  - . 3 x LED (APP, Power, Status)
- 1~2x internal SATA HDD
- Node management with IPMI V2.0 Compliant
- NEBS Level 3 Design: FCC/CE/UL

**Planned 1H'17  
EA Q1'17**


# Switch Node MXN-3610 and extension card

## Specifications

- ½ width in 19" rack
- Intel RRC FM10840 (960Mpps)
  - 4x PCIe x8 to 4x CPU Module
  - Up to 360G for network interface
    - 20x10G on board;
    - 16x 10G from I/O extension card, or to backplane
    - Optional 1U front panel with 40G/100G for OCP
- COMe Type 6 for CPP
  - Intel Core i7/i5/i3 v4 Platform or Atom Platform
  - 1x RJ45 management port for CPP remote management
  - 1x Console port for CMM and CPP local debug
  - 2 x mSATA for redundant storage
- Chassis Management (CMM), IPMI 2.0
  - 1x RJ45 management port for CMM remote management
  - Inventory Info collection (Product Serial, Mfg Date,...)
  - Remote Reset/Power up/down, Adaptive Fan
  - Reset CPP via CMM

## Highlight

PCIe x32 Gen3 Connectors to CPU modules & 200G interconnection to another switch module –  
Total BW ~ 456 Gbps



**Available Now**

# Switch Node


## MXN-3620 w/ MR-IOV support

### Specifications

- ½ width in 19" rack 1U/2U depending on needs
- PLX (Broadcom) PEX87xx/97xx
  - 10x PCIe x 8 to 10x CPU module
  - Up to 160G for network interface
    - 4x 40G QSFP+ on face plate
- COMe Type 7 for CPP
  - Intel Xeon Processor D
  - 1x RJ45 management port for CPP remote management
  - 1x Console port for CMM and CPP local debug
  - 2 x mSATA for redundant storage
  - 1 x4 PCIe to another switch module
- Chassis Management (CMM), IPMI 2.0
  - 1x RJ45 management port for CMM remote management
  - Inventory Info collection (Product Serial, Mfg Date,...)
  - Remote Reset/Power up/down, Adaptive Fan
  - Reset CPP via CMM

### Highlight

PCIe x8 Gen3 Connectors up to 10 CPU modules & x4 Gen3 interconnection to another switch module –  
Total BW ~ 640 Gbps



**Available 1H'17**




# I/O Nodes



Signal Connector

I/O



<b>CSA-Z4X01</b>	NIM1: 4-port GbE copper with LAN bypass
<b>CSA-Z5C4F</b>	NIM2: 4-port GbE SFP w/o LAN bypass
<b>CSA-Z8X10</b>	NIM3: 8-port GbE copper with LAN bypass
<b>CSA-Z5C2F</b>	NIM4: 2-port SFP+ without LAN bypass
<b>CSA-Z5C8F</b>	NIM5: 8-port GbE SFP w/o LAN bypass
<b>CSA-Z5C8F+</b>	NIM5: 8-port GbE SFP+ w/o LAN bypass

## Highlight

PCIe x8 Gen3 Connectors for 1x IO modules – total BW ~ 64 Gbps

**Available Now**

Building Forward Together

# Proven Software Infrastructure

