Edge Computing - Infrastructure for Low-Latency AI/Media

ADLINK Technologies
Jeff Sharpe, Director Product Strategy &
Business Development

Standards Leadership & Technology Partners

Developing, Innovating, and Implementing for the Industry

PICMG Standards

SGeT Standards

PICMG
Executive Member

Intel® IoT Solutions Alliance
Premier Member

OPEN PLATINUM

Quadro Embedded Jetson Ecosystem
Partner

AXIe
Consortium Strategic Member

ONAP
OPEN NETWORK AUTOMATION PLATFORM

Open Network Foundation
Innovation Contributor Member

Edge Computing
Consortium Member

ETSI MEC
Committee Member

Future Airborne Capability Environment
Consortium Member

Wireless Innovation Forum
Member

Central Office Re-architected as a Datacenter
Innovation Contributor Member

OpenEdge Computing
Consortium Member

AMD
Premier Embedded Partner

PXISA
Sponsor Member

Object Management Group
Member

SGeT
Founding Member

NXP
Partner

Wind River Hardware/Software
Gold Partner

OpenFog
Consortium Member

PC/104 Embedded
Consortium Affiliate Member

Automated Imaging Association
Member

Windows Embedded
Partner

Telecom Infra Project
Member

VITA Standards Organization
Member

Montavista
Partner

Association of Gaming Equipment Manufacturers
Associate Member

Edge Compute for low latency

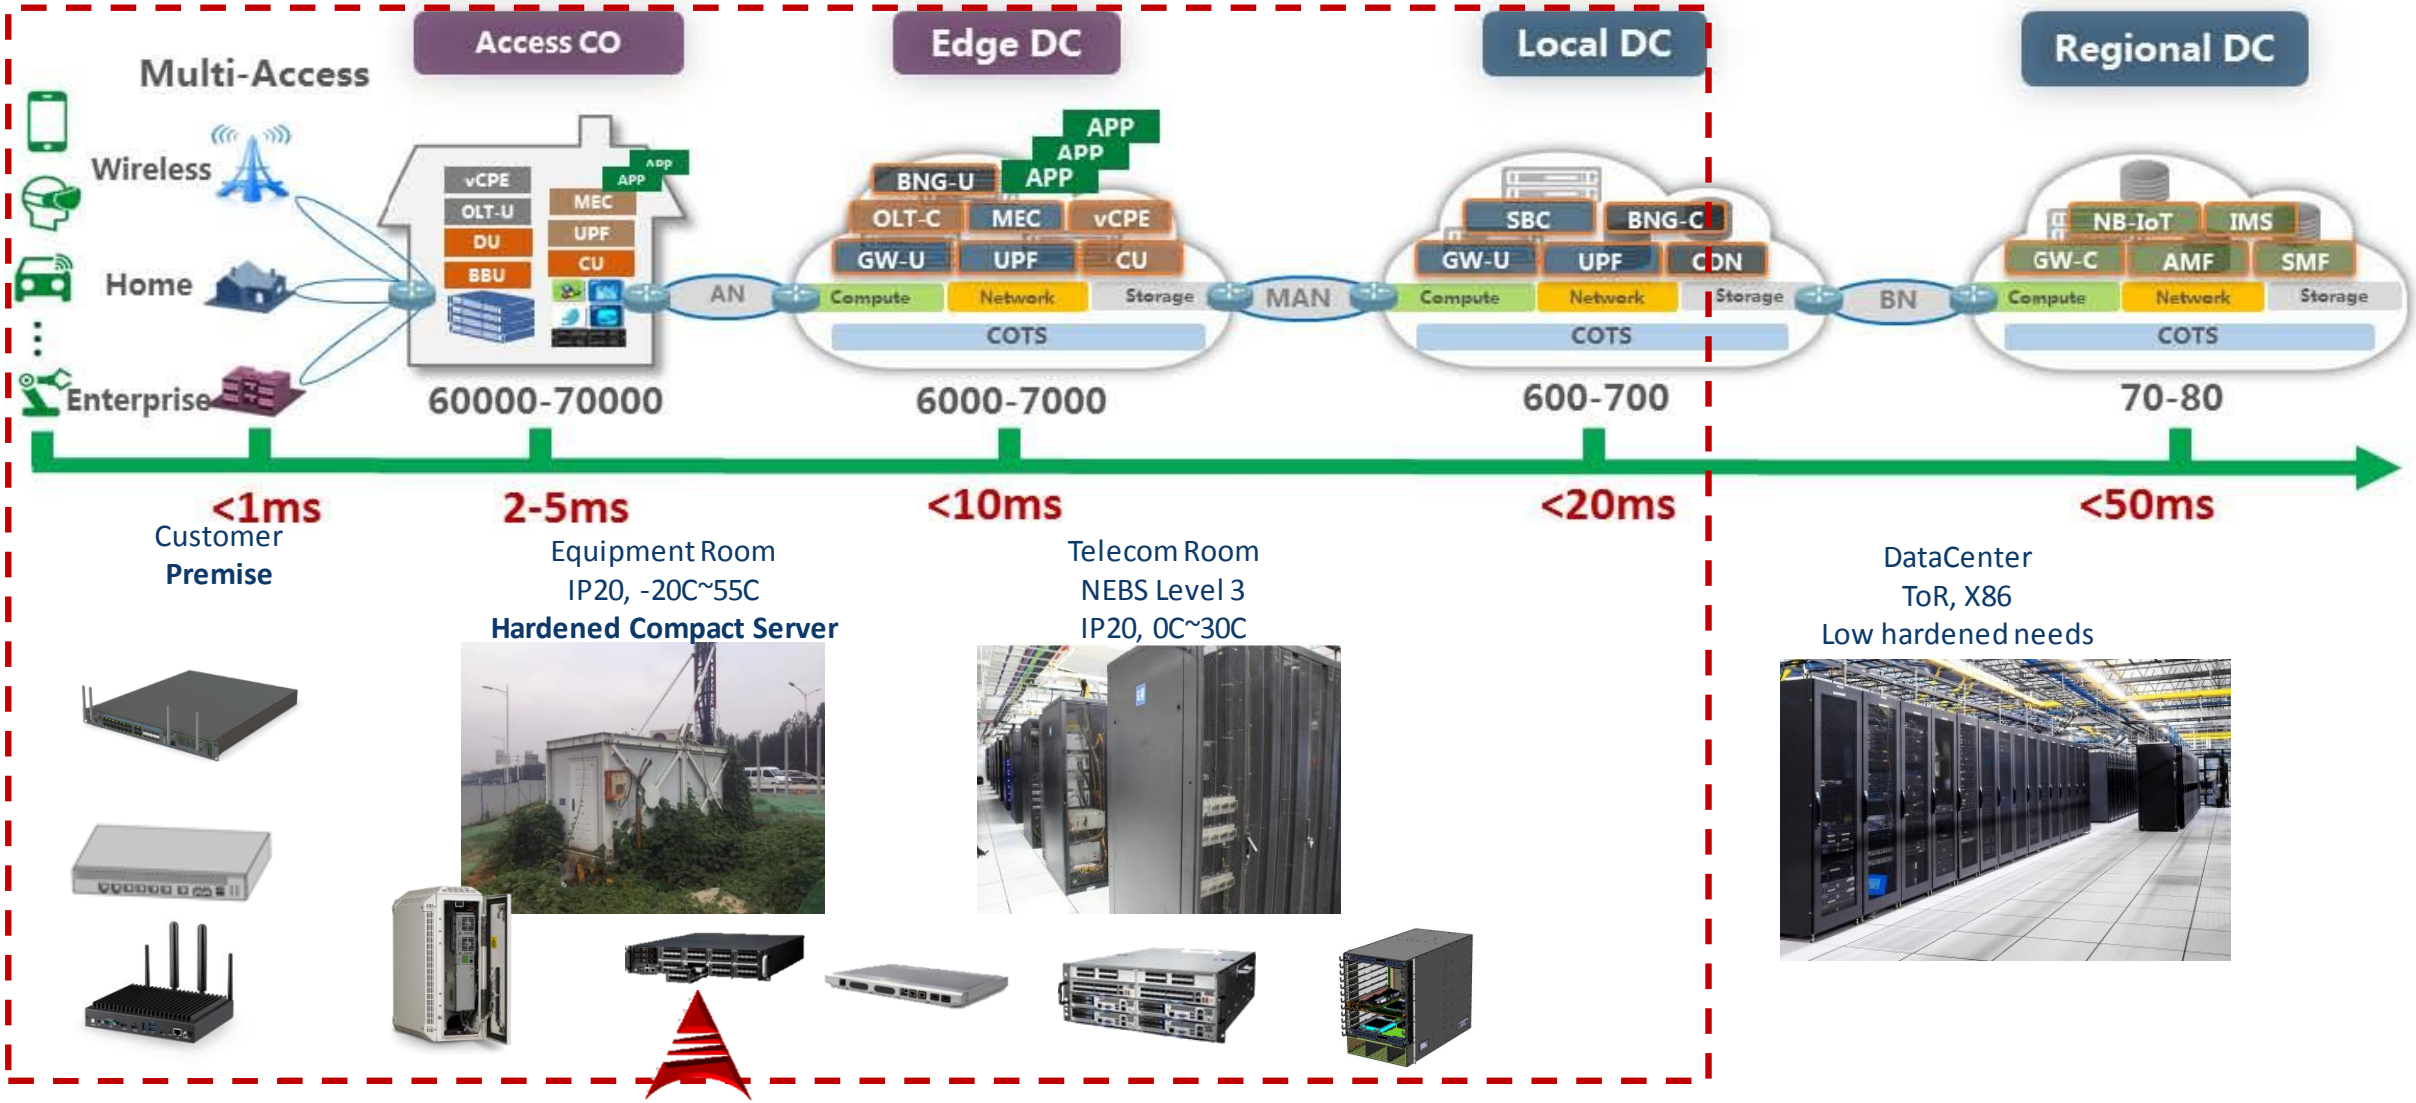

Distributed Edge – Inference / Virtualization

Open Eco-System for RAN deployments (Pole Mount / Container – based Appliances)

Edge Applications

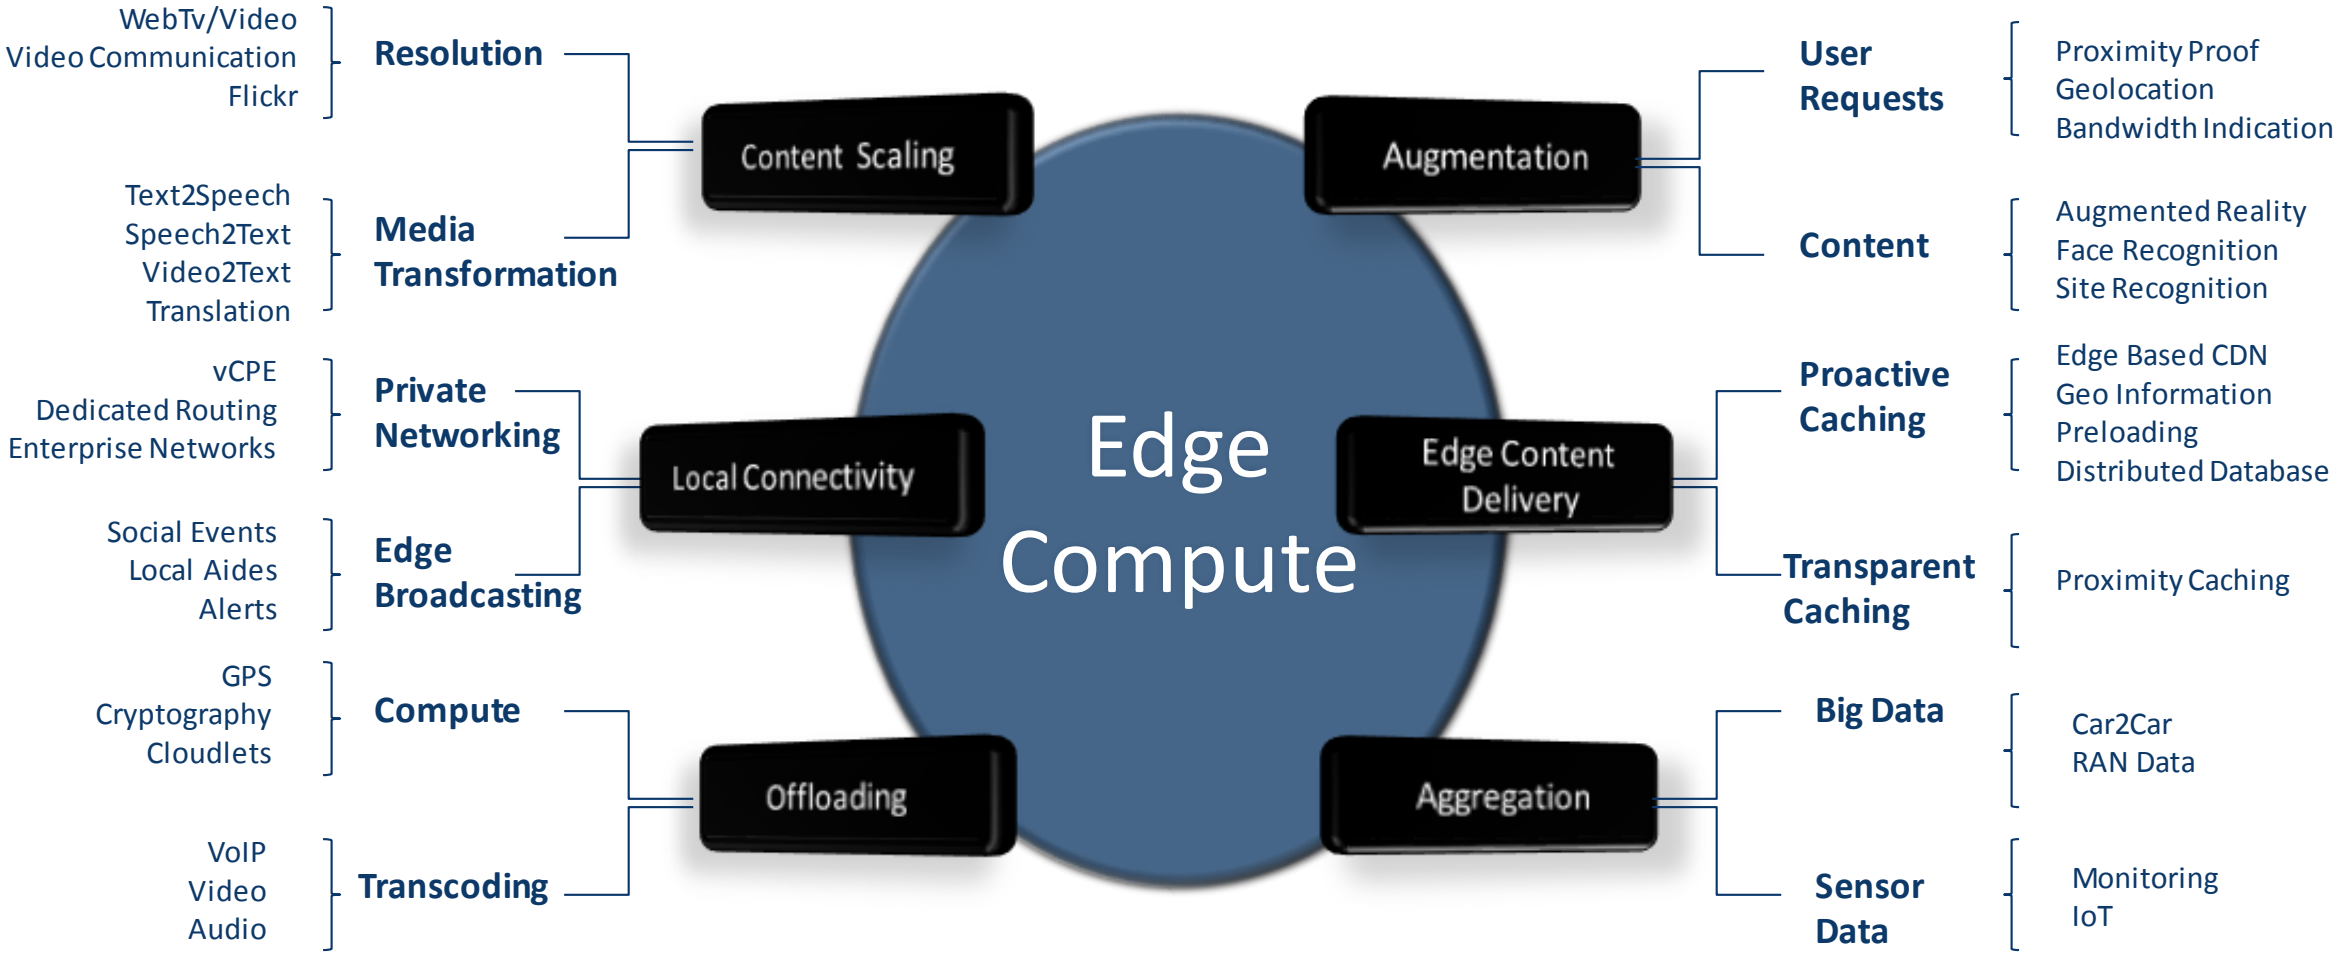

Edge Market Prerequisites – Zero Touch

Vision of Next Generation Networks at the Edge

Self Healing and Resiliency at the Edge
Edge Datacenters are remotely managed and unmanned. 5G provides low latency and high throughput delivery for services – some of which are expected to be up 99.9999% of the time

Automation / AI / ML
The use of GPGPU technology to enhance the responsiveness of issues, policy and real-time analytics. Built in AI for container and vNF management is becoming more acute to Telecoms

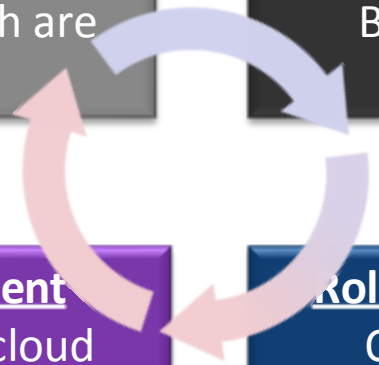

Cloud native approach and DevOps enablement
Service Providers & NFV vendors are driving a cloud native approach. Service providers acknowledge the evolvement of software in every part of the telecom network due to SDN and NFV allowing cloud-based automation for infrastructure deployment, operations and VNF management.

Role of Open Software and Hardware Architectures
Open Architectures for software and hardware infrastructure to develop, deploy and manage vNFs, containers and APIs for services – enablement of mulitiple, integrated partners
ONAP, ETSI, Akraino, ONF OCP TIP and others providing guidance and standards for delivery.

ADLINK & Partners to ensure low risk deployments with technology

Open Architectures – Leaders/Members

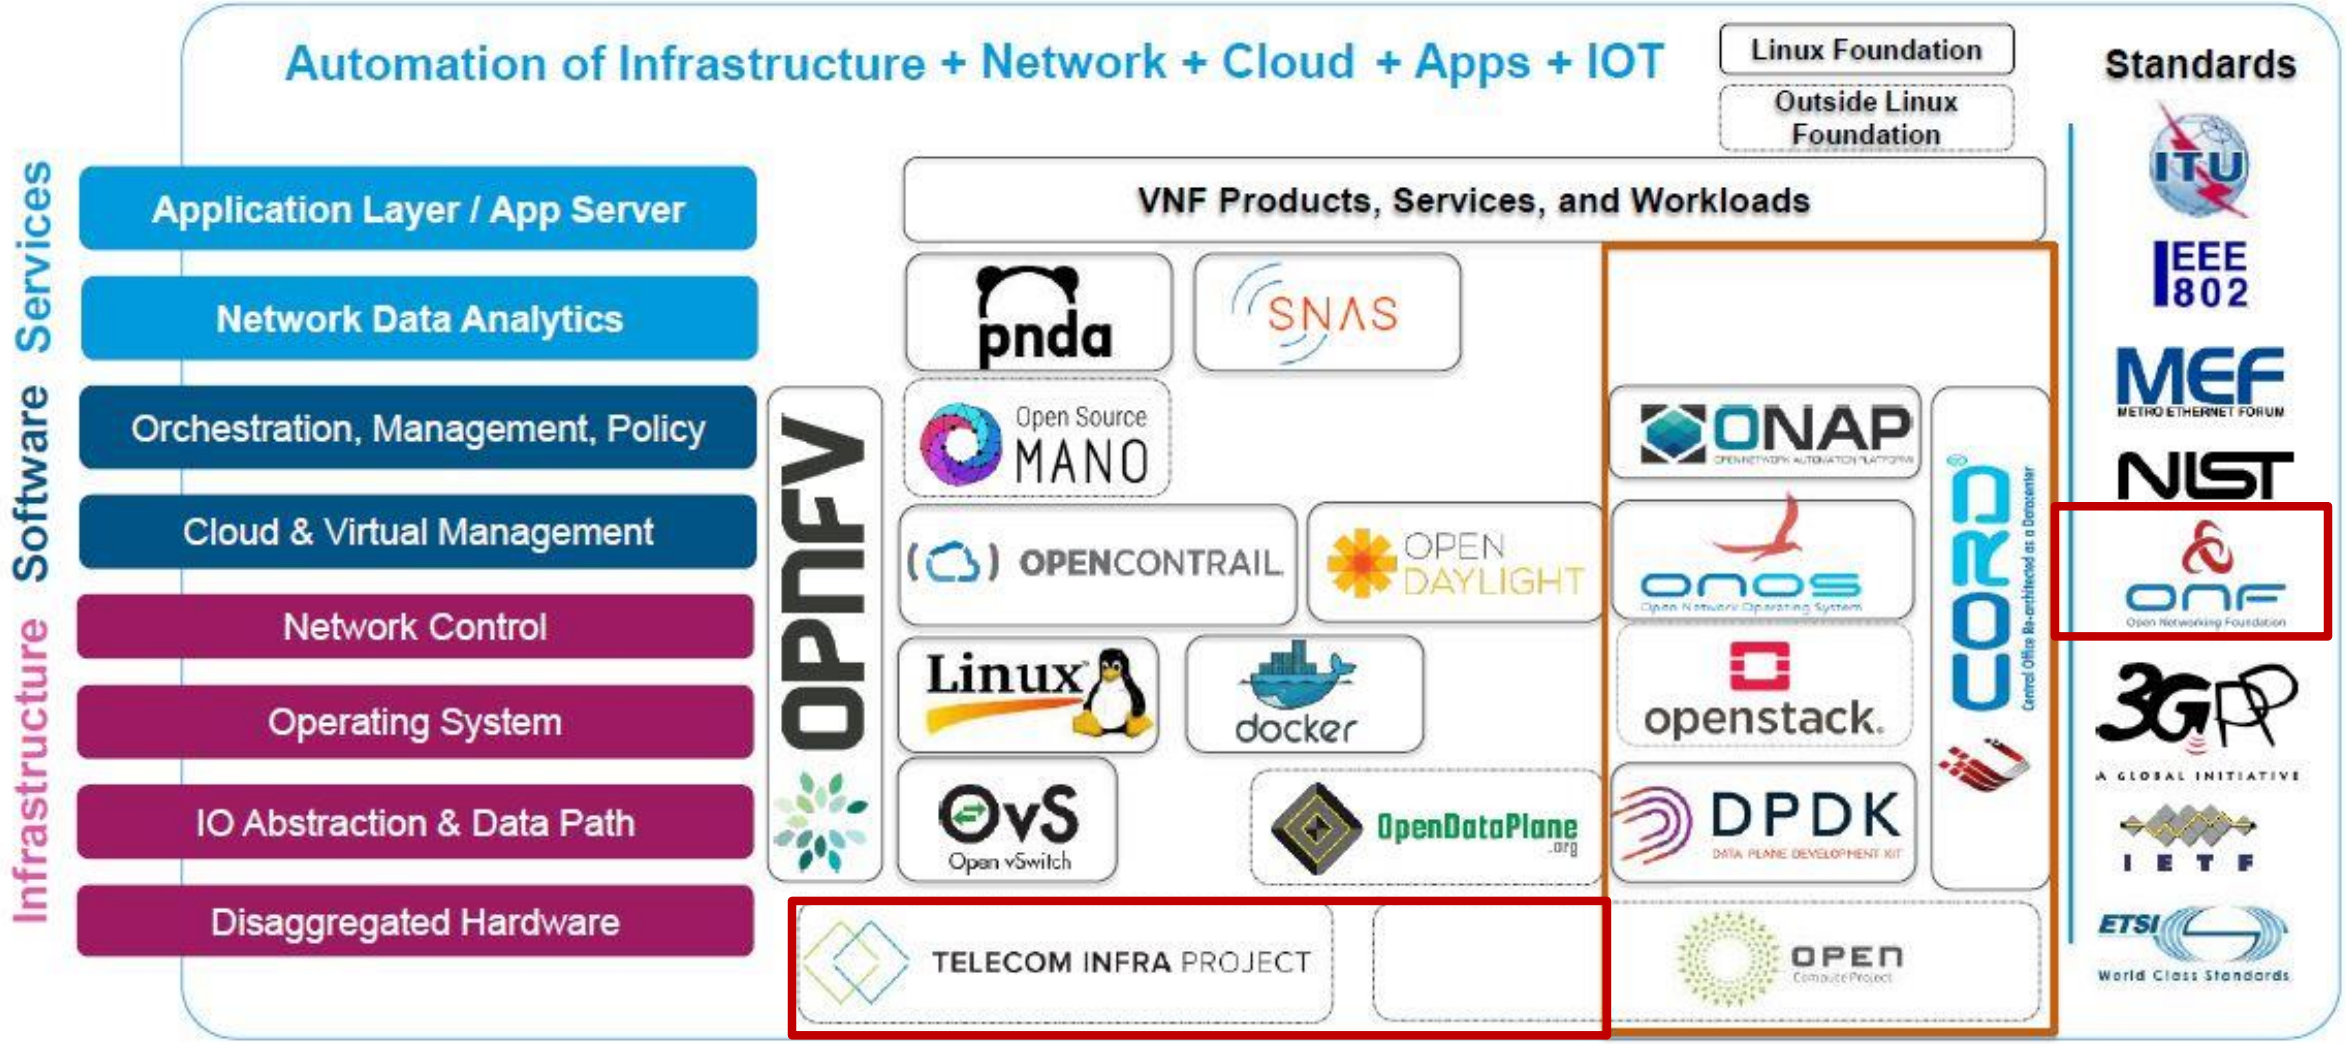

Network Edge Portfolio

Common Architecture

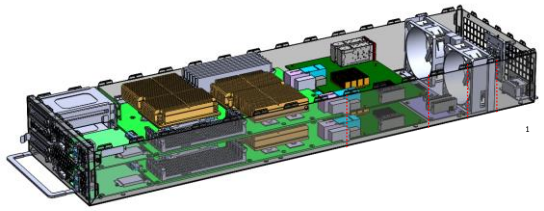

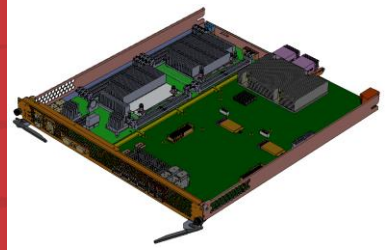

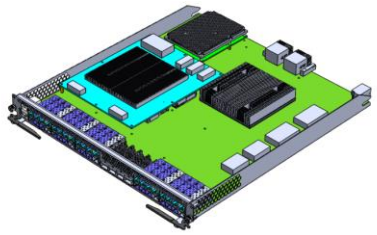

CSA-Z4X01

CSA-Z5C4F

Adlink VPU card

NVidia GPU card

Intel/Xilinx FPGA

Edge Data Center

Backbone & IDC

Metro

Access

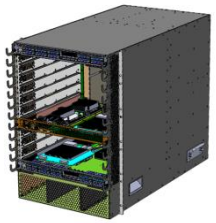

OCP-CG Rack Core

- Mix & Max Compute & Storage Sleds
- X2 Management Switches
- X2 Data Switches
- Balanced Solution:
  - 8x CPU (16 sleds) = 64 Sockets
  - 16x Storage sleds

CSA-7600 Orthogonal System

- P1: New develop Switch Card with 8x100G+40x10G;
- P2: 4xE5+RRC as new Service card for double performance/density
- Full width slieds – reusable assets

CSA-74xx/72xx System

- Switching up to 360Gbs
- Mix/Match Sleds
- NEBs Ready
- Support of GPU and HW acceleration
- Multiple power options
- Reusable assets

NVIDIA Partnership – AMEC Platform

Accelerated MEC – Development Kit for Telecom

AI / DL

Higher Performance Compute

GPU Virtualization

Analytics

Accelerated MEC is an enhanced version of Edge Computing using GPU architectures to support, high performance computing at the edge.

			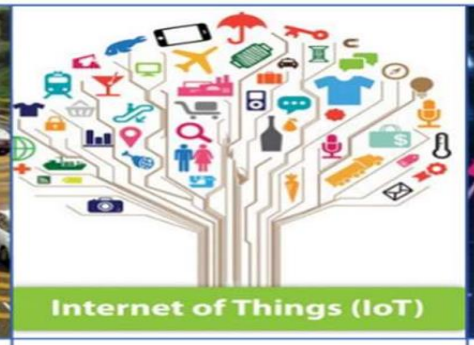	
To provide gigabit broadband connectivity to residential properties as a last mile complement or replacement for fixed networks	To build countrywide mobile experiences that enable new use cases driven by augmented reality/virtual reality (AR/VR), ultra high-definition (UHD) video, artificial intelligence and so on	To deliver reliable, low latency networking to mission-critical businesses to boost efficiency and productivity	To develop industrial ecosystems around machine-to-machine (M2M) and the internet of things (IoT)	To deliver nationwide infrastructure as a service (IaaS)

Phased Approach to AI/ML

Tech Introduction | Market Feedback | Build open architecture standard

High

Performance / Density

Low

ALPS 4800
4U 19” Carrier grade
8 GPU’s P4/P40/V100
2 Xeon Scalable Family
High I/O

ALPS 2200
2U 19” Carrier grade
2 GPU’s P4/P40/V100
2 Xeon Scalable Family
Low-Med I/O

AMEC 4800
Packet Manager
GPU Hypervisor
CUDA toolkit/SDK
NFV/SDN software

Rugged Enclosure
Reuse of 2U system in Outdoor Enclosure
Used for Carrier For small cell deployments

NVIDIA
Embedded Sled
With and w/o NVLink

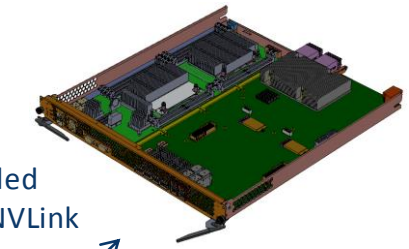

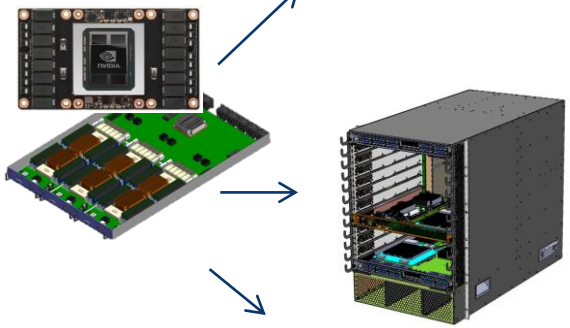

Hardware Development Kit –> Software Infrastructure -> Standard Product

2018 | 1H’19 | 2H’19

Phase I – Dev Kit w/ Infrastructure Software

4U 8 GPU

2U 2 GPU

2U 2GPU RAN

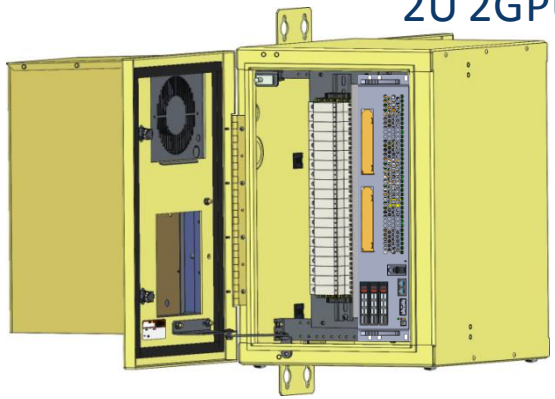

NVIDIA DGX-1 Software Stack

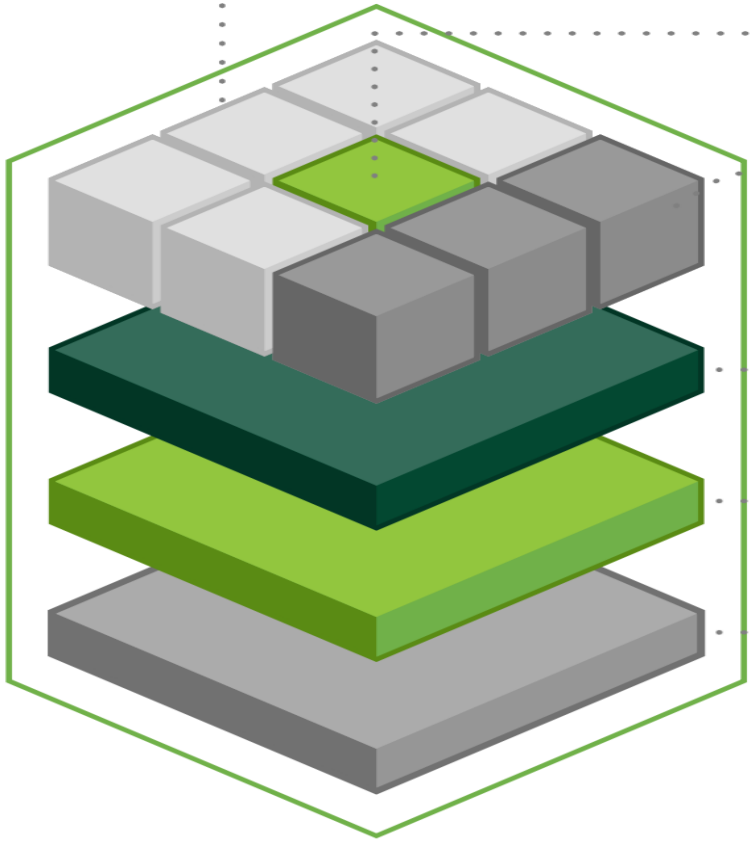

DEEP LEARNING FRAMEWORKS

Caffe · Caffe2 · mxnet · TensorFlow · theano · torch · PYTORCH

DEEP LEARNING USER SOFTWARE

NVIDIA DIGITS™

THIRD PARTY ACCELERATED SOLUTIONS

MAPD · kinetica · Blazing · graphistry

CONTAINERIZATION TOOL

NVIDIA Docker

GPU DRIVER

NVIDIA Driver

SYSTEM

Host OS

MECS-7210 Specification

Reusable CPU asset, introduction of Xeon-D Type 7

MECS-7210 Highlights	
Compact Dimension	Height (2U) x Width (19”) x Depth (420mm)
Expansion	Two dual-slot full-size PCIEx16 3.0 slot; Support GTX-1080Ti & Tesla P100/V100
Acceleration	Up to 100G QAT for symmetric encryption and authentication, asymmetric encryption, digital signatures, RSA, DH, and ECC, and lossless data compression
Extreme Performance	Fully use total 12 memory channels in Skylake GEN

Platform	
Form Factor	**2U 19” with 420mm Depth**
Processor	Intel Xeon Processor Scalable Family (Purley, Cascade Lake compatible)
CPU Socket	2x LGA3647
Chipset	Intel C622~C627 (Support up to 100G QAT)
Memory	16xDDR4, Up to 512GB (RDIMM)
Storage	
On-board SSD	1xCFast
SATA bay	Up to 3x2.5” drive bays*, hot-swappable
Expansion	
PCIe Slot	2x single-slot full-size PCIe 3.0 x16 slots with external panel Or 2x dual-slot full-size PCIe 3.0 x16 slots, internal
IO	
Ethernet	Up to 4x10G SFP+
Console Port	1xRJ45
USB3.0	2x USB on Front plane, 2x Internal USB dock on IO board
Power/Reset	1x Power button, 1x Reset button

MECS-7210 Highlights

- Configuration for expansion platform
  - ✓ (Option #1)2x single-slot full-size PCIe 3.0 x16 slots
  - ✓ (Option #2)2x dual-slot full-size PCIe 3.0 x16 slots

- 3x 2.5” SATA bay
- RJ-45 console port
- 2x RJ-45 10/100/1000BASE-T Ethernet ports
- 2x 10G SFP+ Ethernet ports (Optional 4x10G SFP+)
- 2x USB 3.0

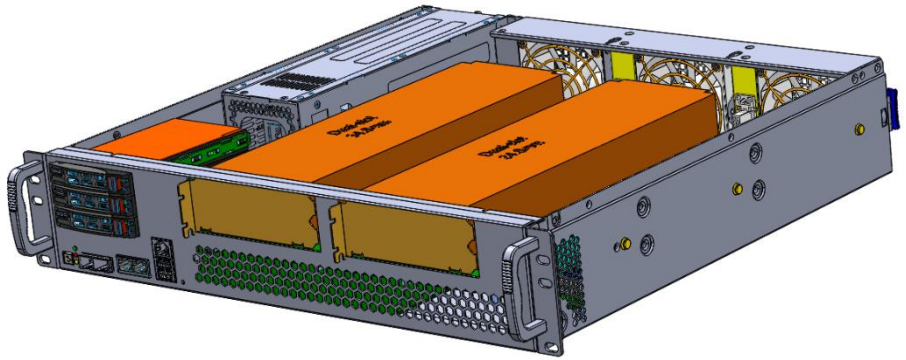

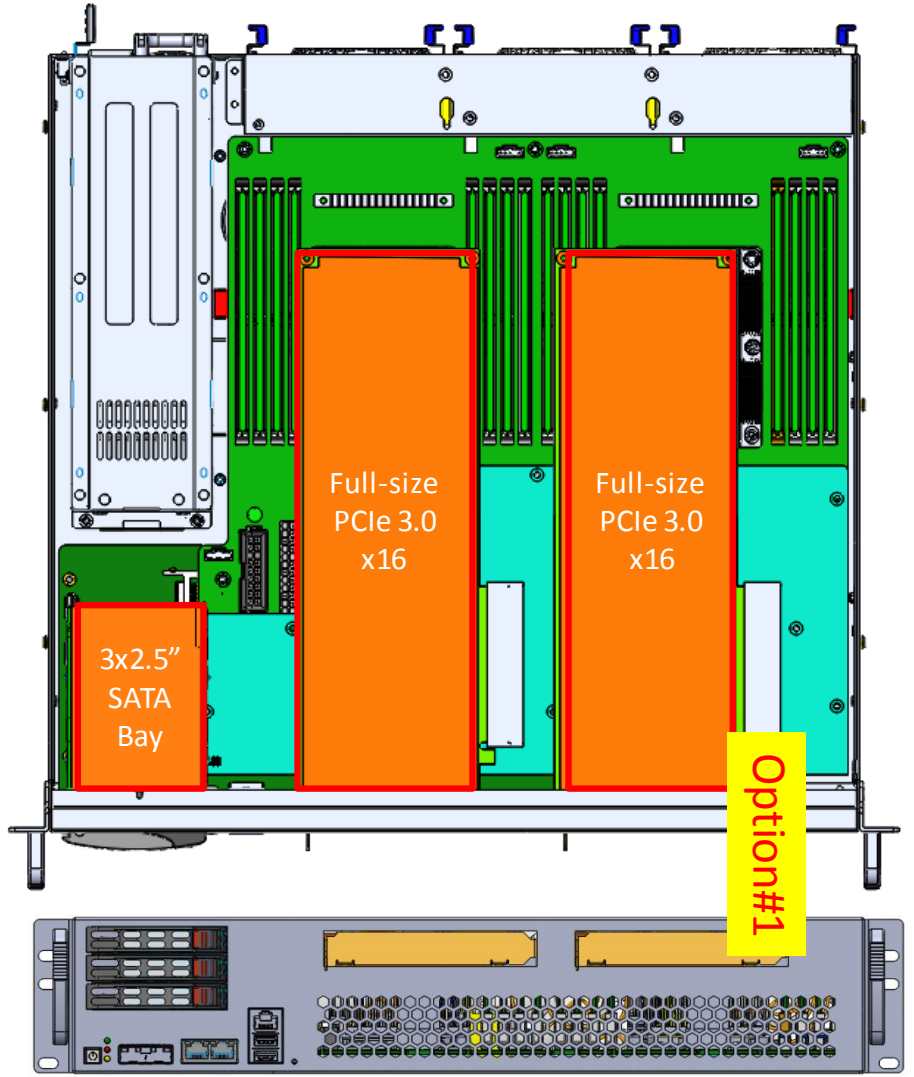

MECS-7210 Block Diagram

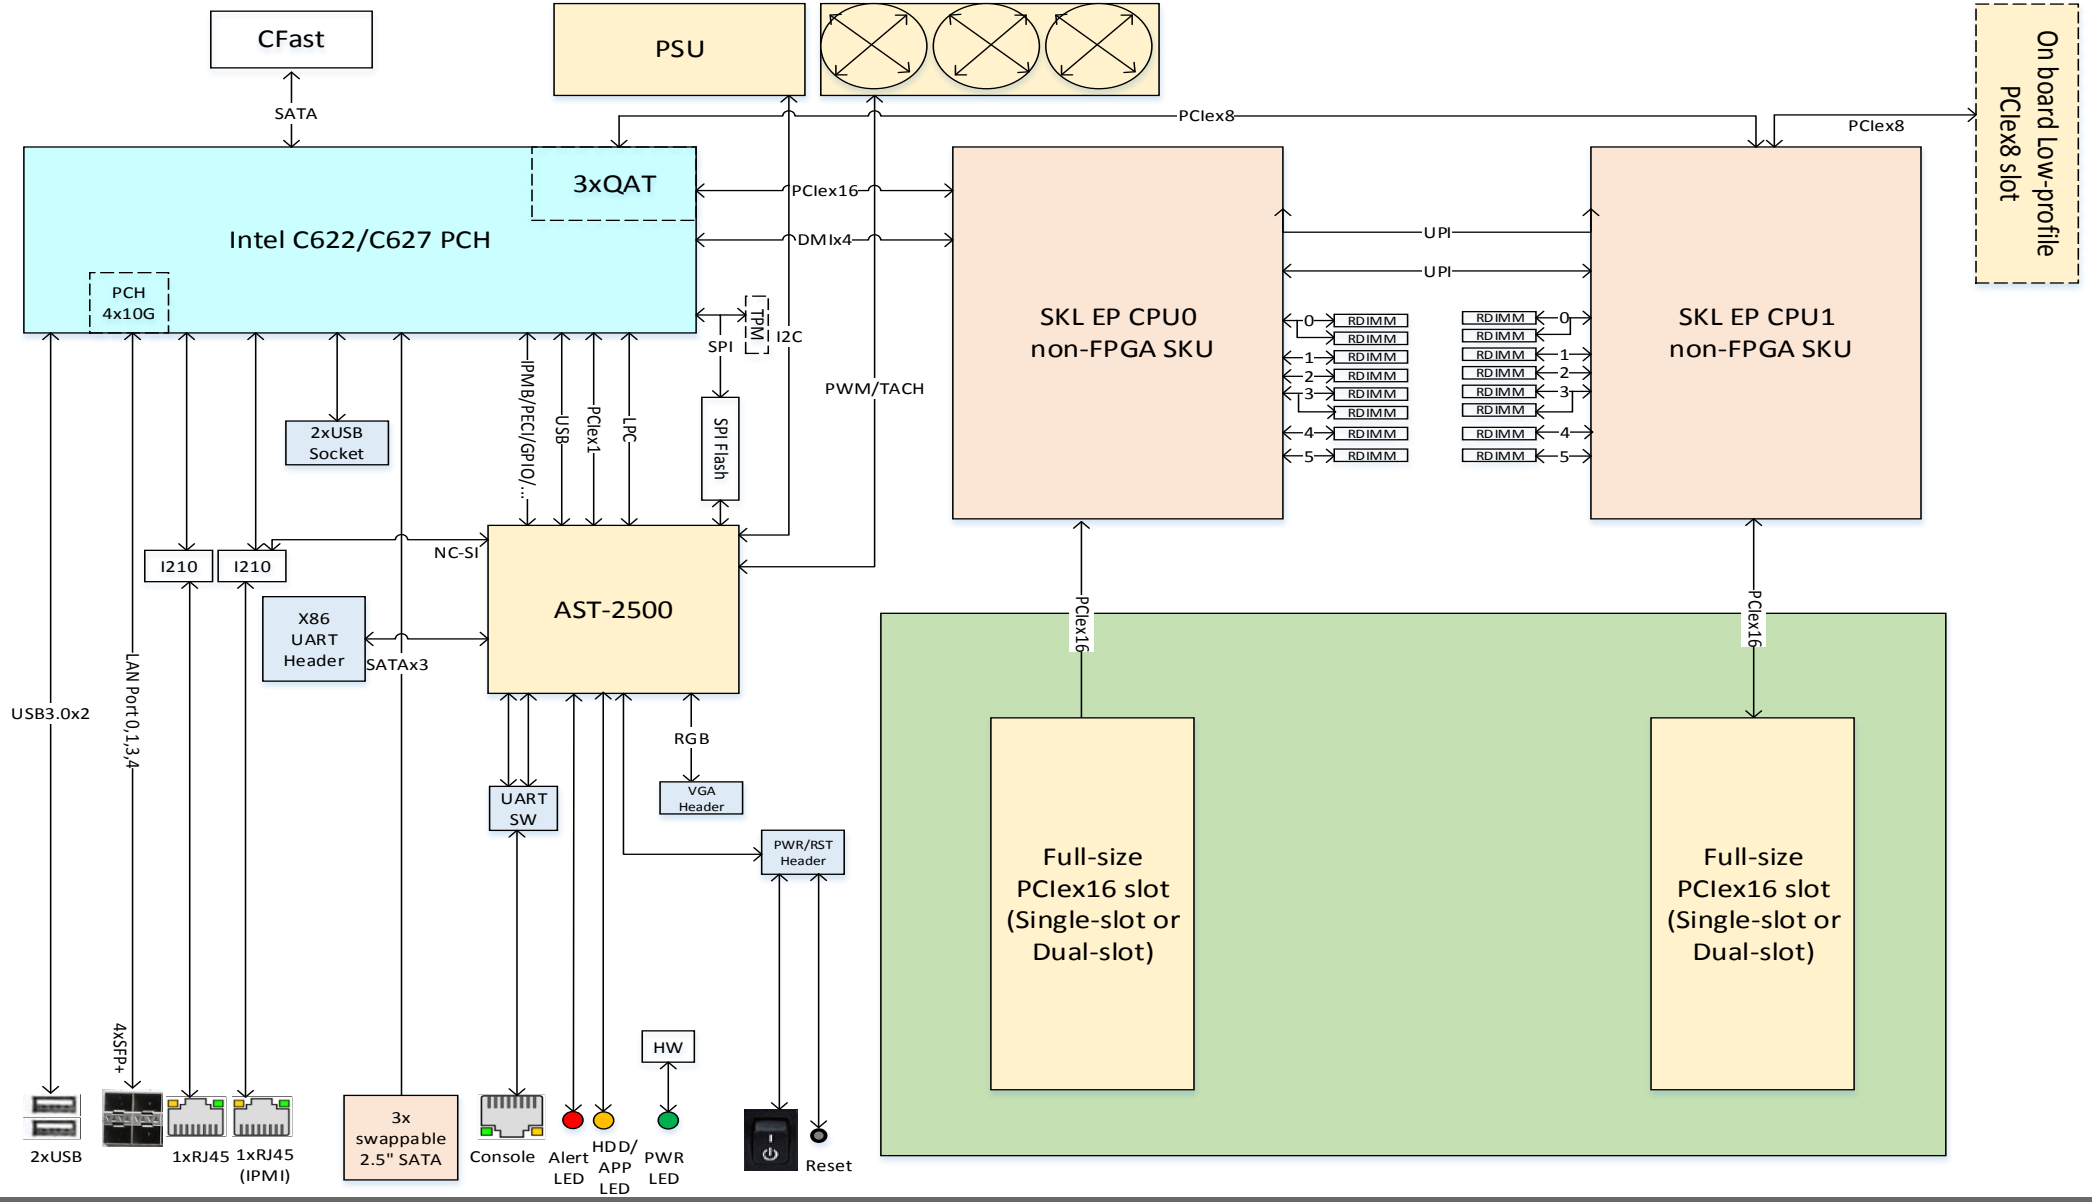

2U MEC Edge Server appliance

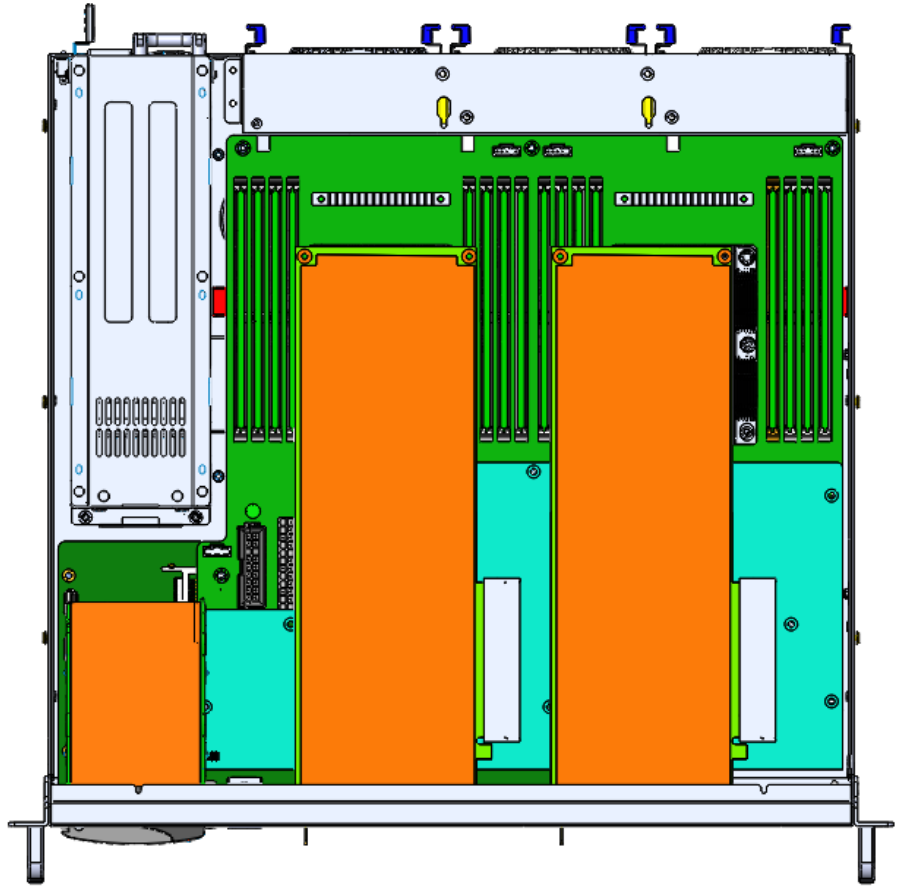

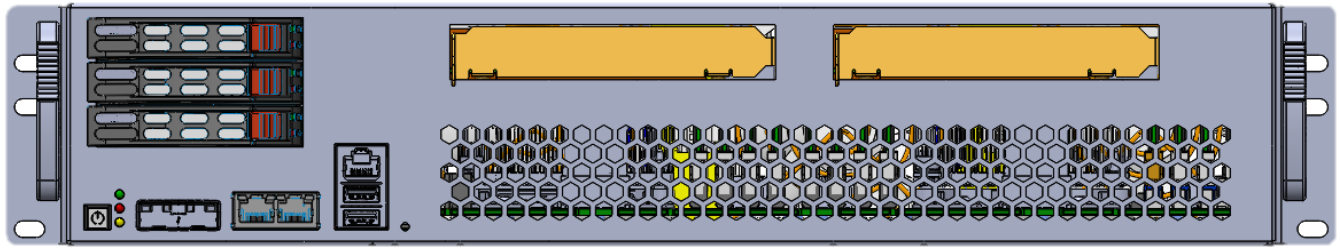

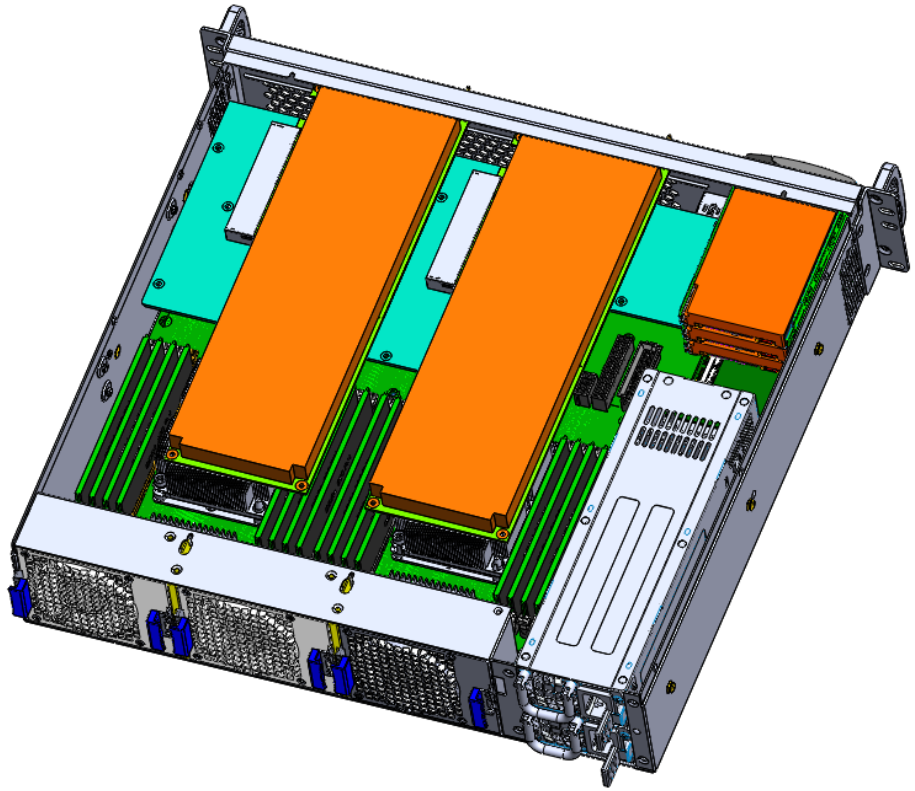

Supported Dual-slot GPU Card in MECS-7210

Supported Dual-slot GPU card	Sample
LeadTek GeForce Series with cooling FAN. Requirement: Thickness <=35mm	LeadTek WinFast_GTX_1080_Ti(10780) 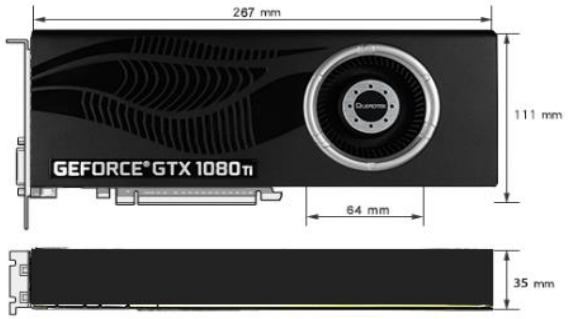
NVIDA Tesla Series with passive cooling . Requirement: Thickness <=40mm	Tesla V100, T4 or P40 280*125*40MM 

Pole Mountable for 5G deployments

4U Pole CUBE with Battery Compartment

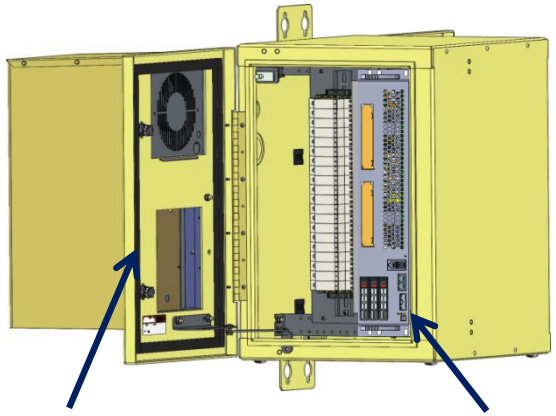

7U Pole CUBE with A/C

- Low Profile Pole Mount Enclosure
  - 4 RU: 24/48” x 11” x 19”
  - 7 RU: 24/48” x 15” x 19”
- Optional Integrated Power
  - Power and battery or power only
- Optional Battery Compartment
  - (1) String 80AH NiCD batteries
- Thermal Options:
  - 580/750W HX
  - 2K BTU HVAC (7U only)
- GR-487 & UL60950 Certified
- Verizon Approved

Call for discussion and combined input

Define OCP Edge Server use best-in-class, modular components

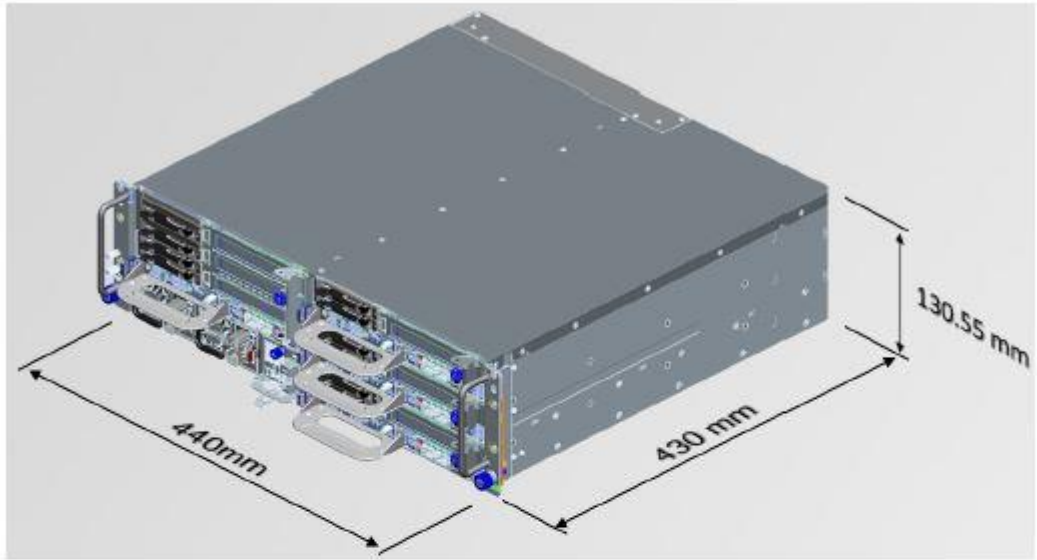

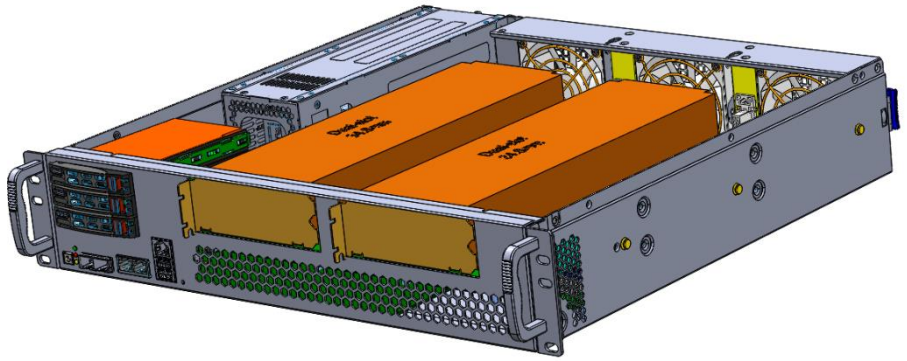

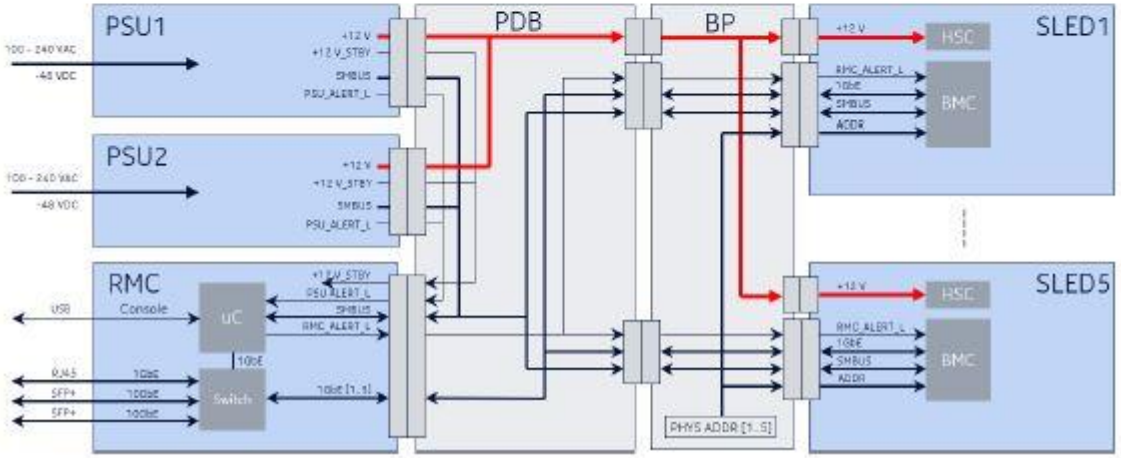

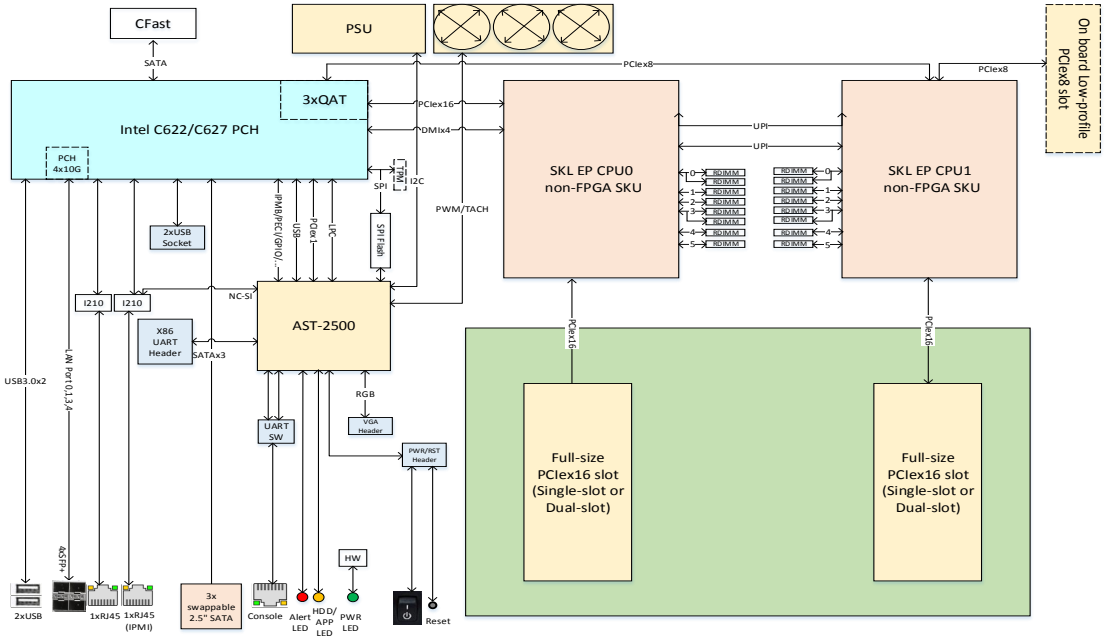

ADLINK

Leading EDGE COMPUTING