




OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.



Advancing Open Architectures to Build Your Server Room

Jeff Sharpe, Director Product Strategy, ADLINK

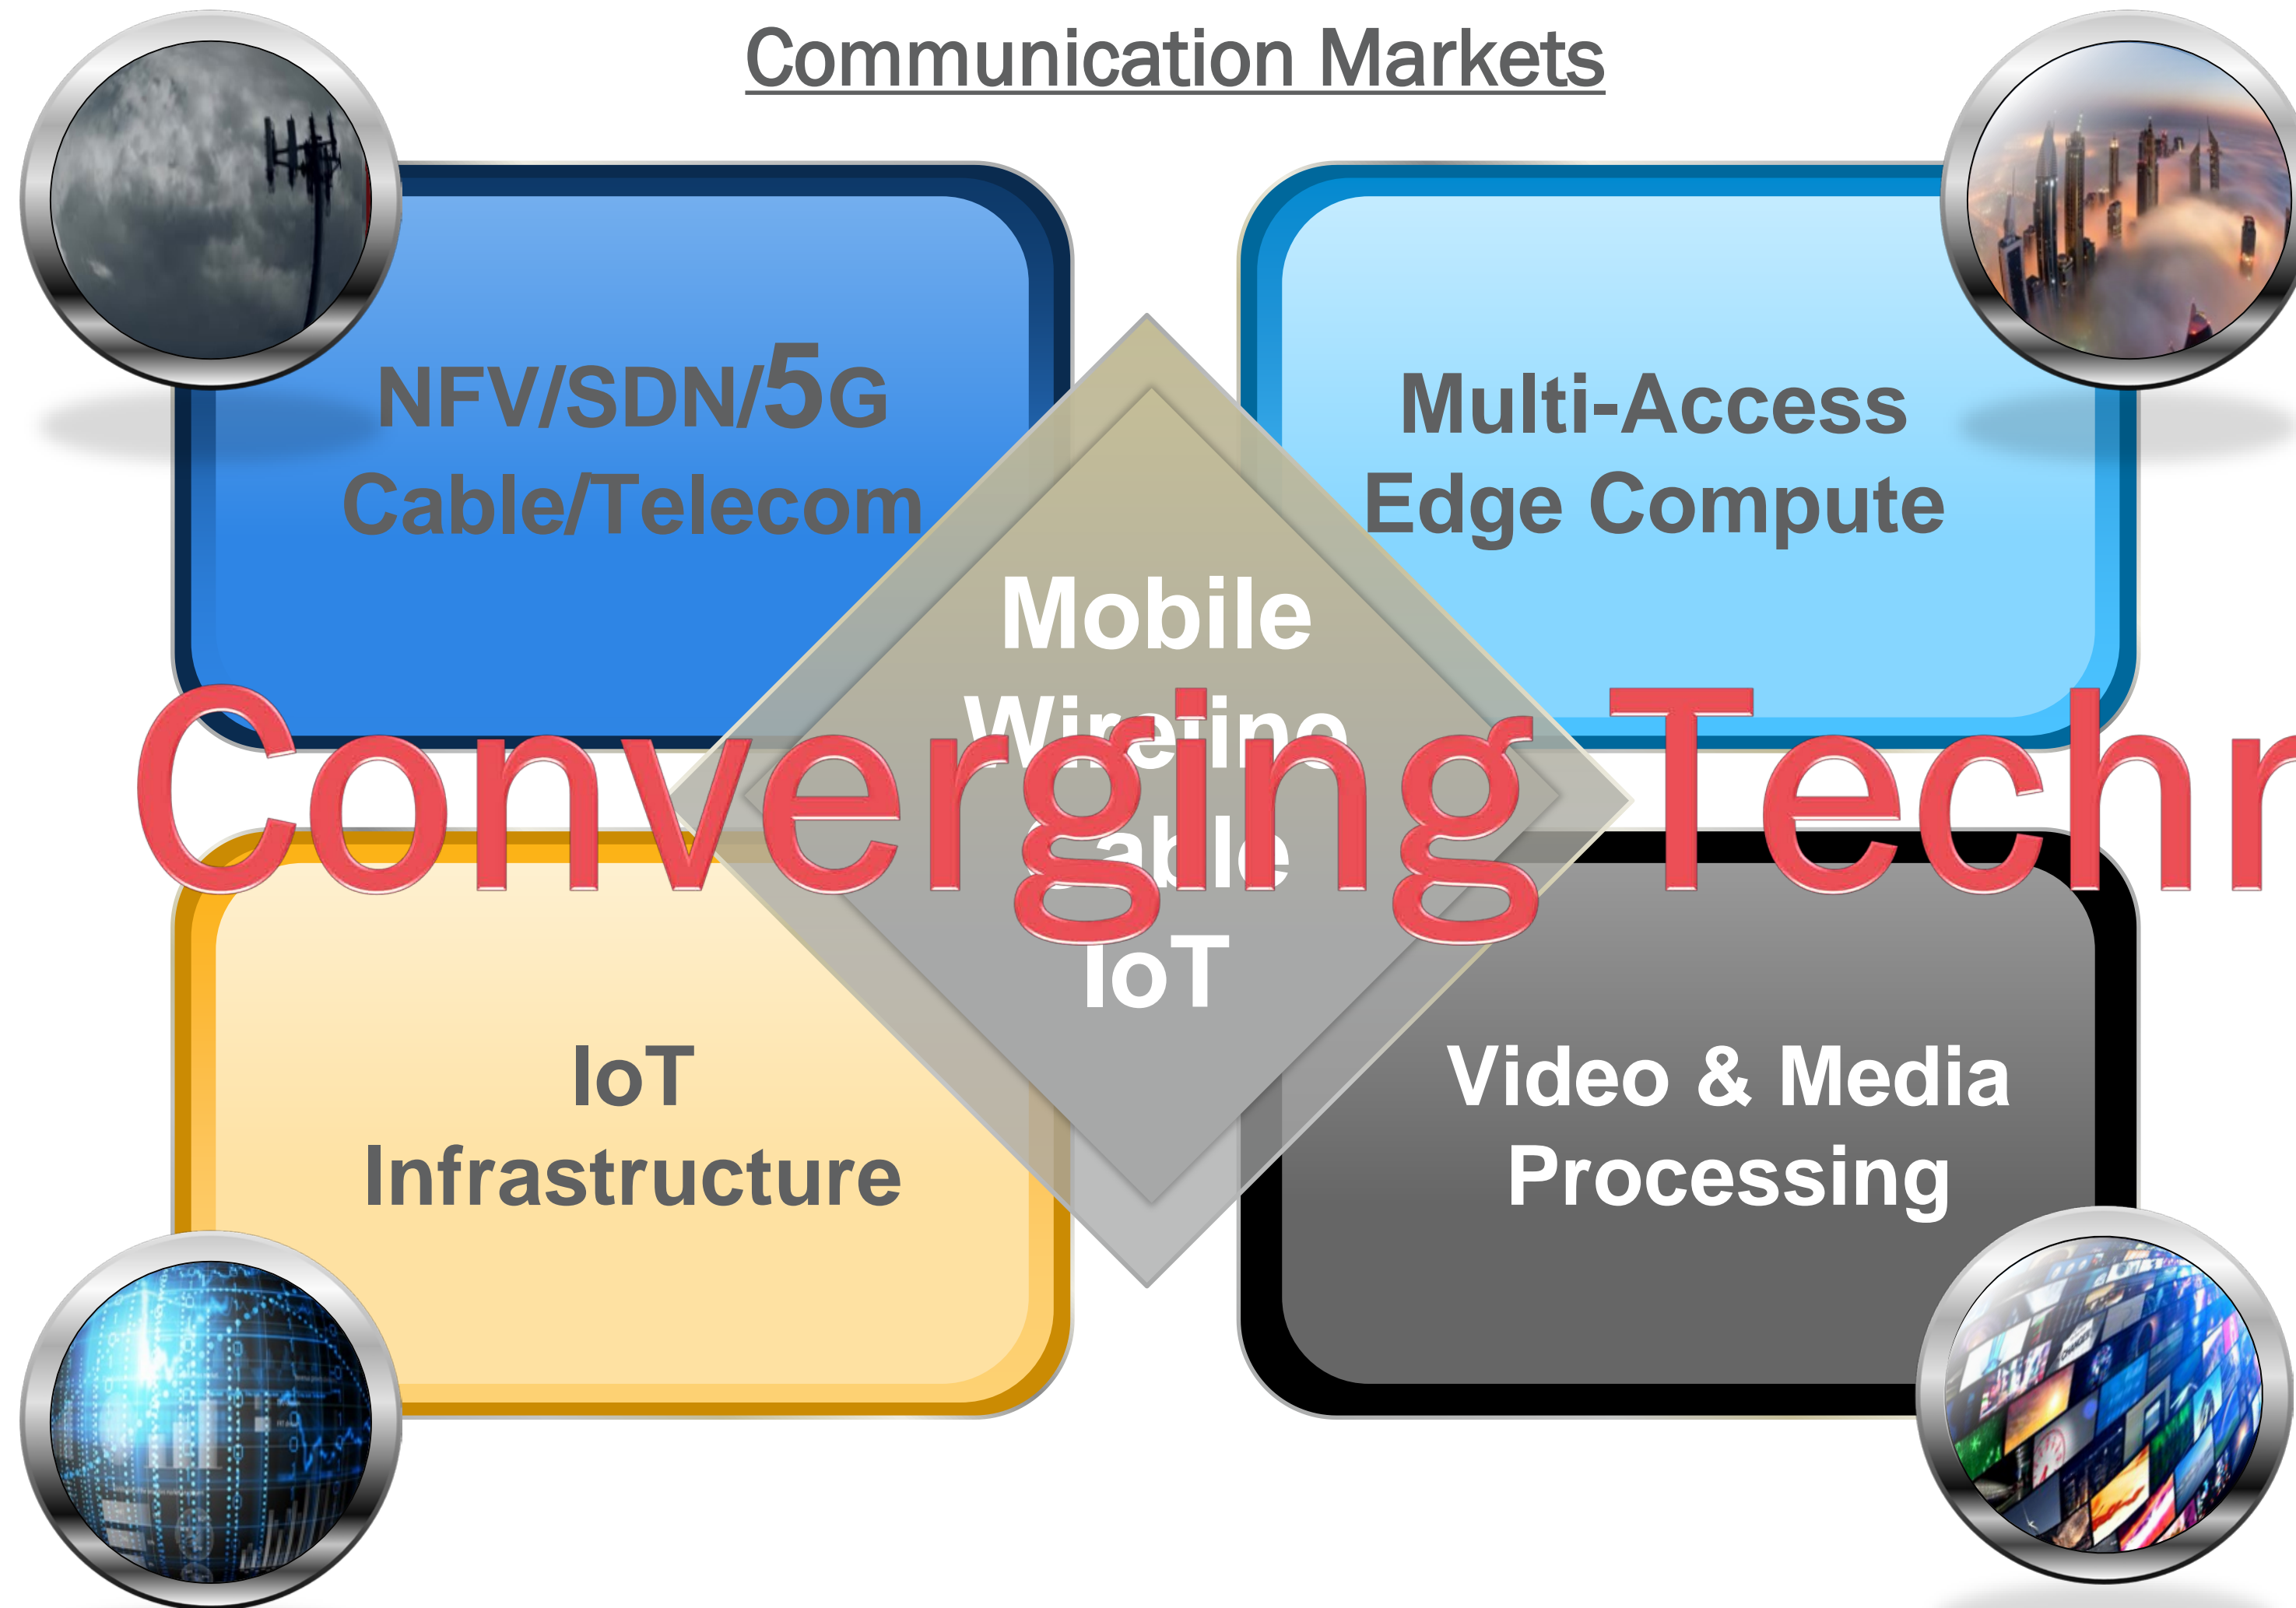
OPEN. FOR BUSINESS.



Converging of Markets

Communication Markets

Adjacency Markets



Defense



**Mining /
Industrial**



Transportation





**Private Cloud
Enterprise**

Needs vs. Value vs. Cost

- Appliance vs. Rack level Architectures?
- Silicon requirements now vs. future?
- Requirements for placement?
- Build-out, Life Cycle Management?
- Vendor Lock-in?
- Proprietary vs. Open?



Open Architectures & Committees



SCORPIO

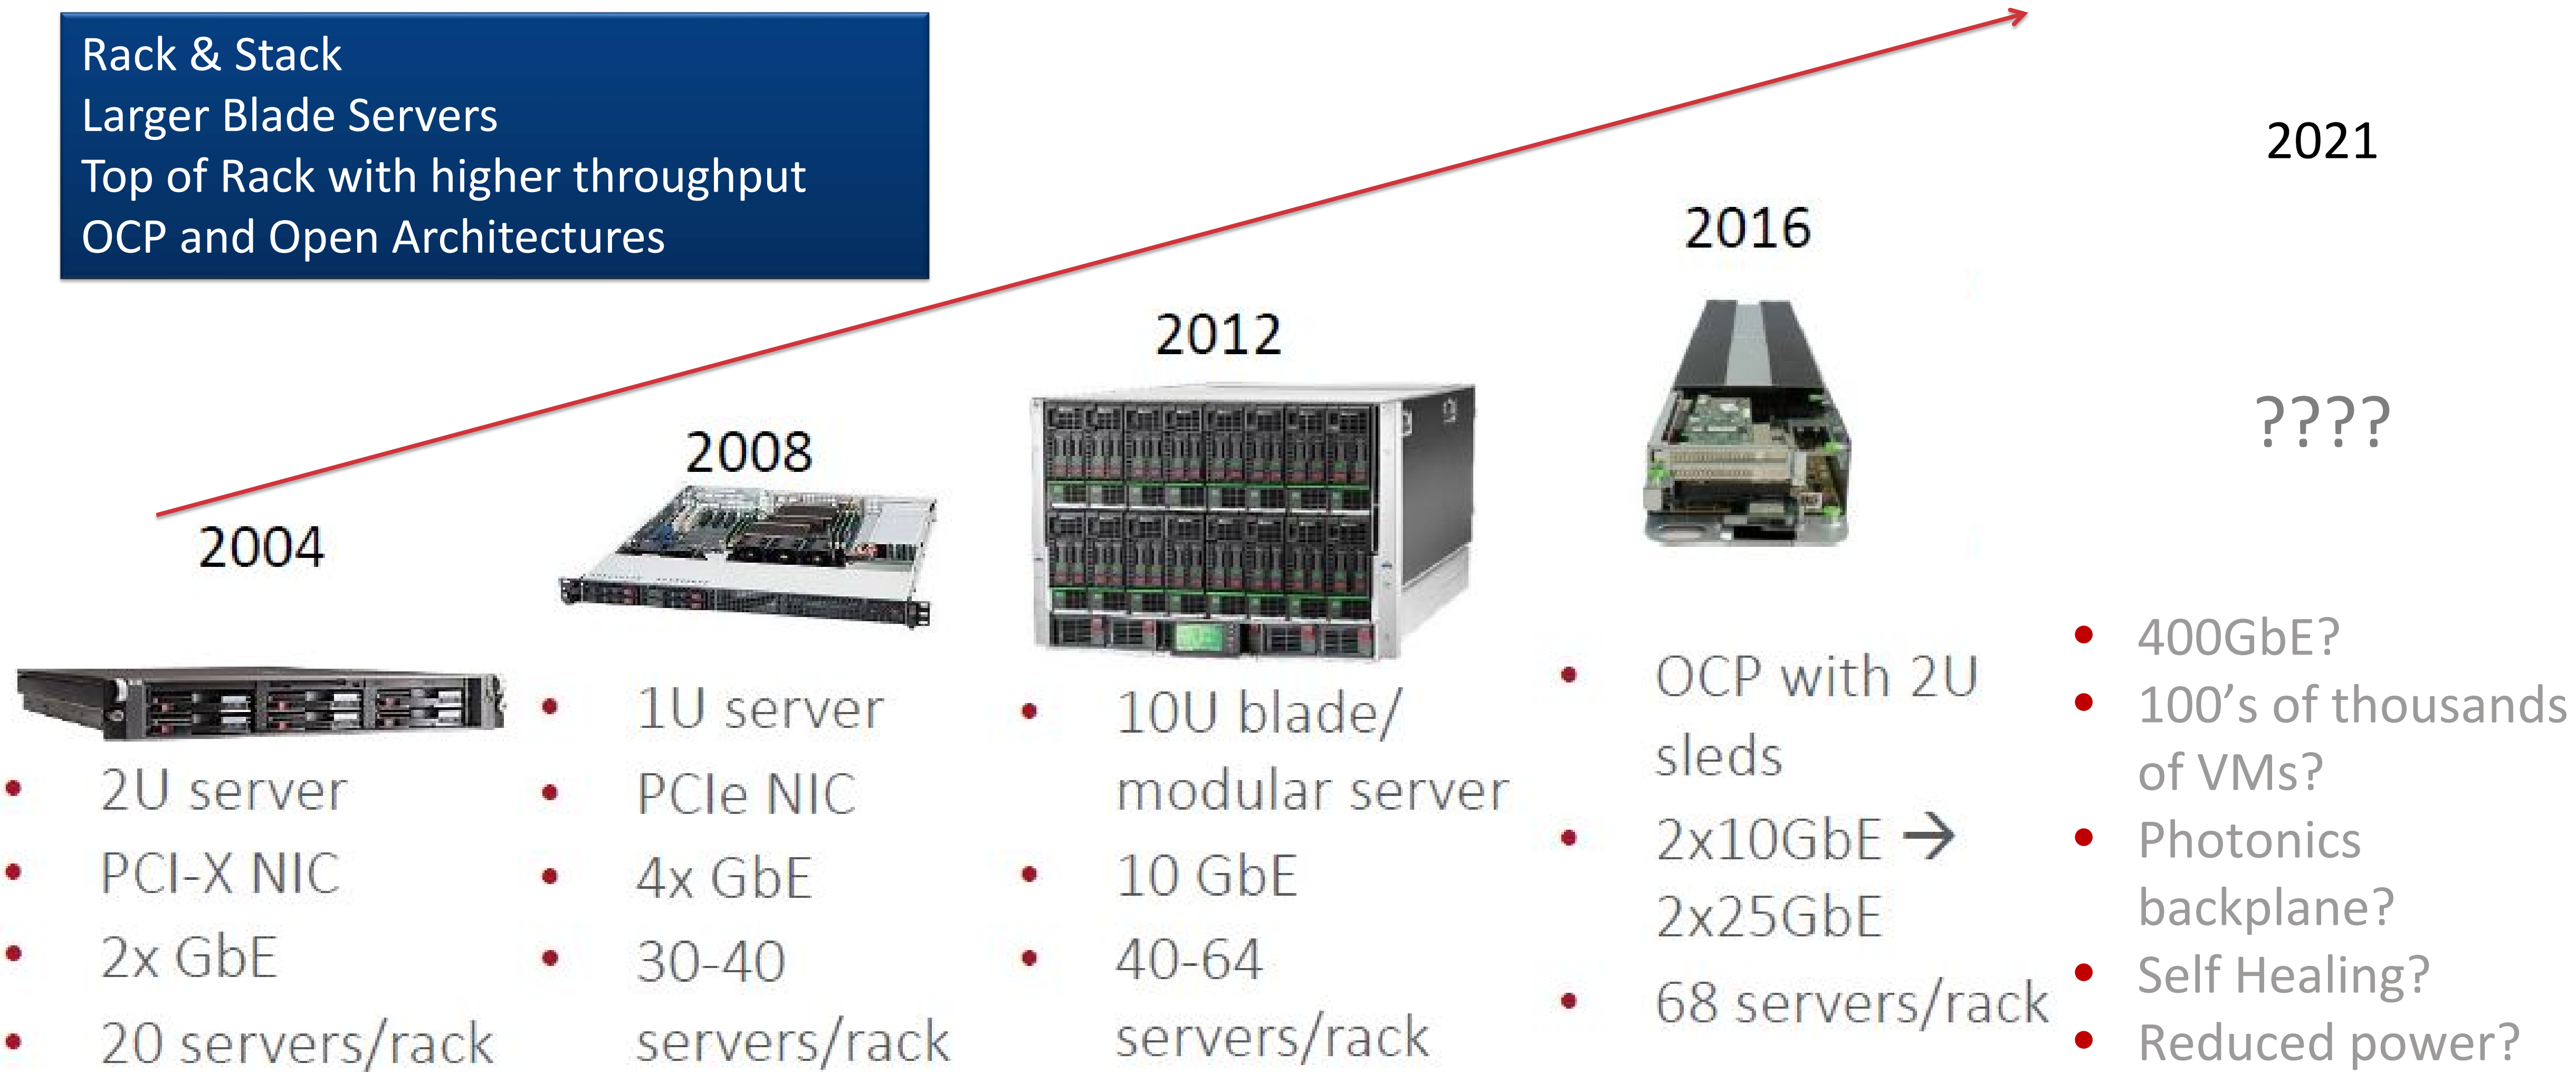


Open Architectures/Committees
Double Edged Sword
What is the right fit?




Looking back – We’ve come a long way

Rack & Stack
Larger Blade Servers
Top of Rack with higher throughput
OCP and Open Architectures



2004



- 2U server
- PCI-X NIC
- 2x GbE
- 20 servers/rack

2008



- 1U server
- PCIe NIC
- 4x GbE
- 30-40 servers/rack

2012



- 10U blade/modular server
- 10 GbE
- 40-64 servers/rack

2016



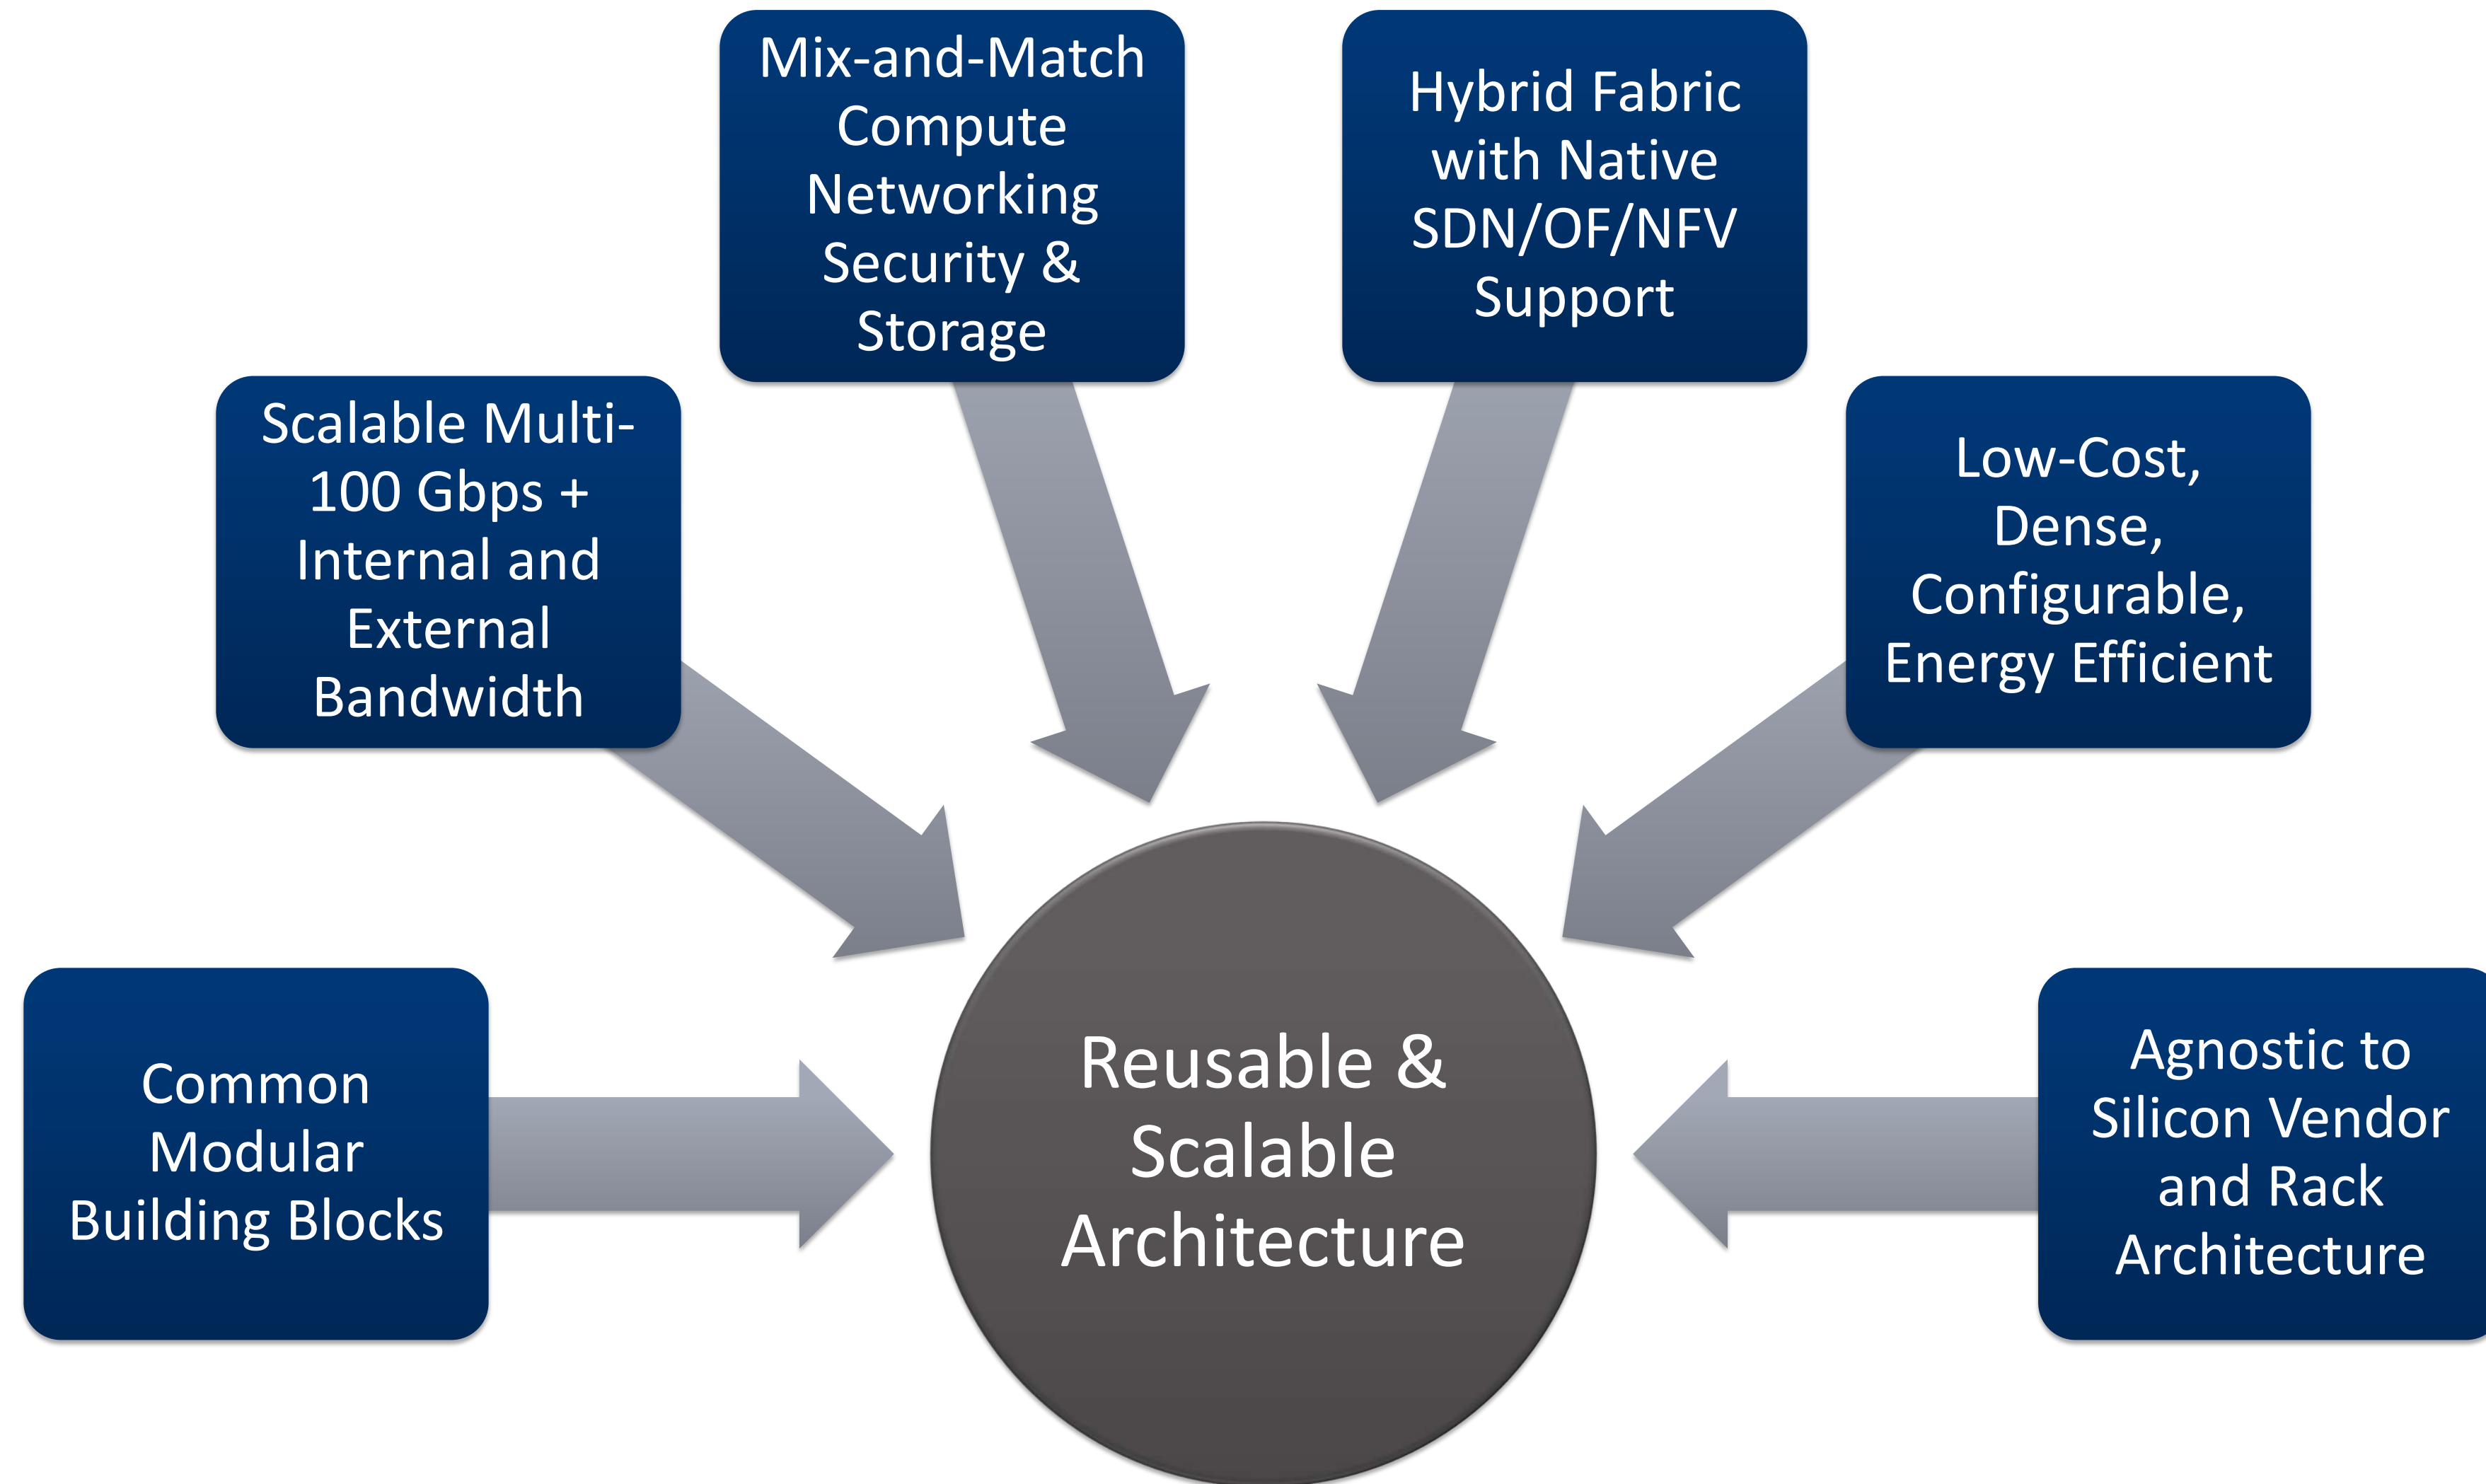
- OCP with 2U sleds
- 2x10GbE → 2x25GbE
- 68 servers/rack

2021

????

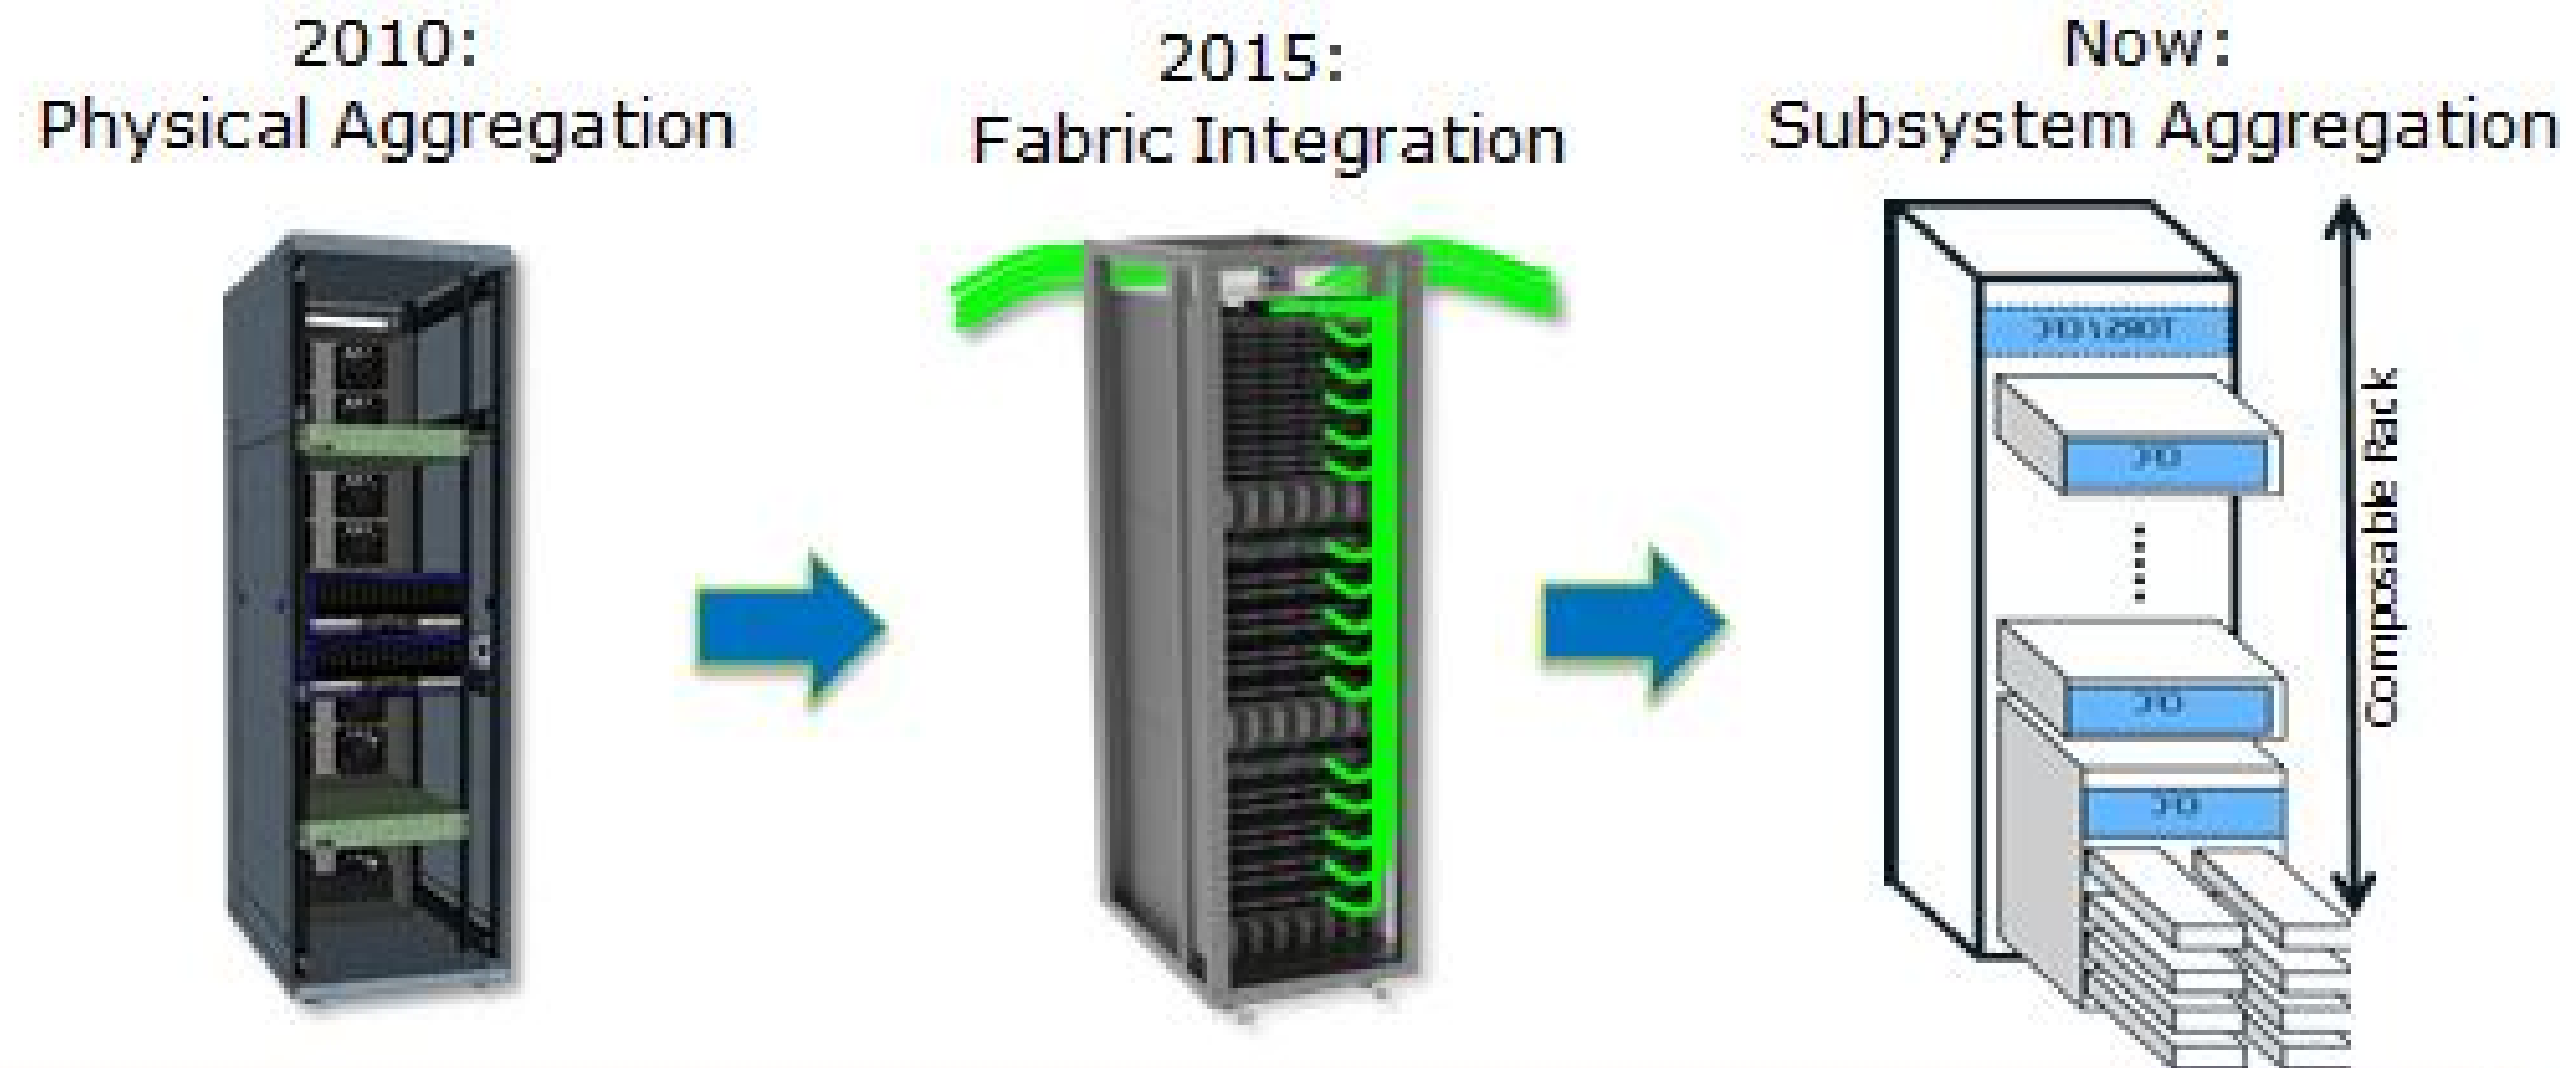
- 400GbE?
- 100's of thousands of VMs?
- Photonics backplane?
- Self Healing?
- Reduced power?

My Checklist -> Future Proofing



RSA -> RSD -> Open architectures

Intel's early vision of RSD evolution to Full Aggregation



Rack level Challenges	<ul style="list-style-type: none"> Power/thermal inefficiencies Limited density 	<ul style="list-style-type: none"> Fixed configuration High cable count Limited scalability 	<ul style="list-style-type: none"> No service based configurability Limited scalability
Intel RSA Solution	<ul style="list-style-type: none"> Shared power Shared cooling 	<ul style="list-style-type: none"> Distributed Switching Flexible topologies Software Defined Networking 	<ul style="list-style-type: none"> Pooled storage/boot Pooled memory Software Defined Server
End-user Benefit	<ul style="list-style-type: none"> Lower TCO 	<ul style="list-style-type: none"> Lower TCO Lower CPU refresh costs Higher compute density 	<ul style="list-style-type: none"> Resources match workload Service scalability

Carrier Grade Platform Evolution

Proprietary -> Open Architectures
Standards driven Platforms (ATCA)
Lower TCO
Large mix of eco-system participants
Utilization of Central Office for Edge
Scale, density and cost effectiveness

OCP OpenRack
Now to future

ATCA / Carrier-Grade Appliances
2001 to Now

Proprietary
1990's to Now



Anatomy of an Open OCP – CG system



Physical

- Suitable for CO retrofit, new telco data center environments & Data Centers
- 19” rack width and standard “RU” spacing for greatest flexibility
- 1000 to 1200mm cabinet depth, supporting GR-3160 floor spacing dimensions

Content/workload

- Heterogeneous compute and storage servers
- Built for SDN and Virtualized systems for optimal performance/ecosystem
- Eco-system for CPU, GPU, ARM, DSP and switching

Management

- Ethernet based OOB management network connecting all nodes via a TOR management switch
- Optional rack level platform manager

Networking/Interconnect

- One or more Ethernet TOR networking switches for I/O aggregation to nodes
- Fiber cables, hot-swappable blind-mate with flexible interconnect mapping.
- Environment, power, seismic & acoustic CO environmental requirements applicable
- Safety and other certification standards also applicable
- NEBS optional (L1/L3)

2017 Evolution of OCP Technology into Telecom



OCP-CG-OpenRack-19"
Frame Level
Specification




OCP-CG-OpenRack-19"
Open Sled Specification



Half width




Full width



Sled Definition based on
OCCERA

Enables additional container for sleds, components and partners for OCCERA growth/future

42U OpenRack Configuration




•42U OCP-CG OpenRack 19”

- 600mm & 800mm wide rack options
- Power → 110/208VAC 3ph & 230/400VAC PDU -48VDC option as well
 - 3 PSU shelves provides 12 x 2500W PSU's
- Management Switches (x2)
 - Switch #1 : Connects 1G to each server BMC
 - Switch #2 : Connects 1G to each server CPU
- Data Switches
 - 1 or 2 switches (up to 3.2 Tbps each)
 - 40G uplinks to spine switch, 10G downlink to each server
 - Option for 100G uplinks & 25G downlinks (v2.3)

•Standard Configurations

- Balanced : 8x Compute (16 sleds) + 8x Storage
- Storage : 17x Storage Shelves
- Same components as 19U frame

Additional 12u frame size at our booth



Typical Rack Commissioning: 1-3 months




- *Schedule contractors*
- *...Receive components*
- *...Schedule different contractors*
- *...Build rack*
- *...Schedule different contractors*
- *....Install & test*
- ...all gated around scheduled maintenance windows*



With OCP-CG delivery: 3 Days

Telco OCP Derived Platforms

Pre-wired & tested rack core, sleds are FRUs with fully optical interconnect to ToR switches



Plug in Power




Connect to Spine




*Slide in sleds
...fast install
...easy serviceability*

Why OCP Open Frame / Open Architectures



- Simplified design
- Simple integrated rack
- Improved Serviceability, Reliability
- High density compute and storage
- Fewer racks, significant real-estate savings
- Integrated pooled power
- Improved redundancy
- Tool-less, wire-free design
- Significant operations savings



OCCERA – Open Architecture for Reuse



-Common Architecture, Extreme Throughput, HW Acceleration



CPU



NIC



GPU/
GPGPU



Specific Appliance card
(H/W accelerator/Security/DPI)



CSA-7200/7210

8 NIMs,
128 Eth Ports

Rich I/O



CSA-7400

4 Compute Nodes,
360G Throughput


High Compute Performance
& Throughput



CSA-7600

12 Compute Nodes
6.4T Throughput

Large Compute,
Performance &
Backplane



OCP

2 Compute Nodes,
NIMS, Storage

Leverages all key
attributes

OCCERA - Network Edge Portfolio


Edge
Data Center



OCP-CG Rack Core

- Mix & Max Compute & Storage Sleds
- X2 Management Switches
- X2 Data Switches
- Balanced Solution:
 - 8x half width sleds 16 sleds 16 sockets
 - Optional full storage frame with 16x Storage sleds
- Roadmap = NVIDIA GPU, Marvell ARM

Backbone
& IDC



CSA-7600
1.4T~2.8T

CSA-7600 Orthogonal System

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

Metro



CSA-7400
400G

CSA-7400/7200 System

- Introduction of next gen Intel Skylake
- Updates to Packetmanager
- RedRock Canyon Switching
- Into of GPU and ARM based sleds
- High throughput and Packet customization
- Mix/Match Sleds

Access



CSA-7200/7210

Use Case - LifeLine DataCenters

Re-Purpose and building a Niche Market for jobs and Secure Clouds



LifeLine DataCenter, Indianapolis, Ind.



Level 4 of 5 FedRamp Certified –
only 3 providers today

Throughput, Ease of Installation, Flexible and Open



Reusable Hardware Configurations:

Config 1: Appliance for smaller customers:

- SKU1: CSA-7400 - 4 CPU Node / 2 Switch Node
- SKU2: AMEC - GPU Edge Dev Kit
- SKU3: CSA-7600 - Orthogonal System

Config 2: OCP Half Rack or Full Rack:

- SKU 1: 100% Storage, 3 Switches
- SKU 2: 60% CPU 40% Storage 3 Switches

Config 3 - GPU

- SKU1: 4U Dev Kit Appliance
- SKU2: GPU OCP Full Rack, Half Rack

Conclusion



- **Open Standards/Architecture**: Lot's to choose from – what works best? We've selected OCP OpenRack for scalability, flexibility and long-term viability
- **Collaboration & Open Architectures**: ADLINK believes in collaboration, open architectures while working with industry leading eco-system partners/committees
- **On-going dedication to OCP CG OpenRack 19"** – Content rich roadmap for converged markets with continued expansion of eco-system providers
- **Extreme Computing**: OCCERA is a base for many types of systems from 1U to 42U focusing on re-usability and scalability using open architectures



OCP SUMMIT