



RapidIO Overview

The Unified Fabric for Performance Critical Computing



Multi-Processor
Embedded Interconnect

Switched | Scalable | Low Latency | Reliable

ANY TOPOLOGY
ANY PROCESSOR
OPEN STANDARD



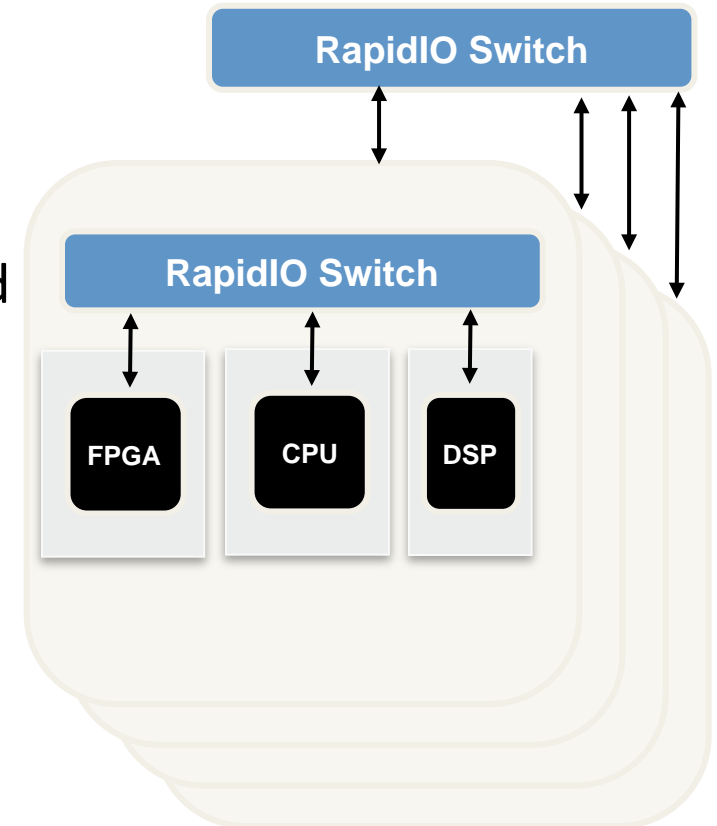
10 Gbps > 20 Gbps > 40 Gbps > 100+ Gbps

WIRELESS INFRASTRUCTURE | SERVER | HPEC | IMAGING | AEROSPACE | INDUSTRIAL

Rick O'Connor – rickoco@RapidIO.org

- RapidIO Overview
 - Dominant Wireless Infrastructure market share
 - ARM 64-bit Coherent Scale Out
- System Challenges
- Data Center Compute & Networking (DCCN)
 - Multi Vendor Collaboration Open Source Platform
- VITA 78 SpaceVPX and NGSIS
- RapidIO based HPC Systems
 - HP Moonshot Proliant m800
 - IDT HPAC Lab using DCCN Hadoop Cluster
 - Open Compute Project HPC effort
- Summary

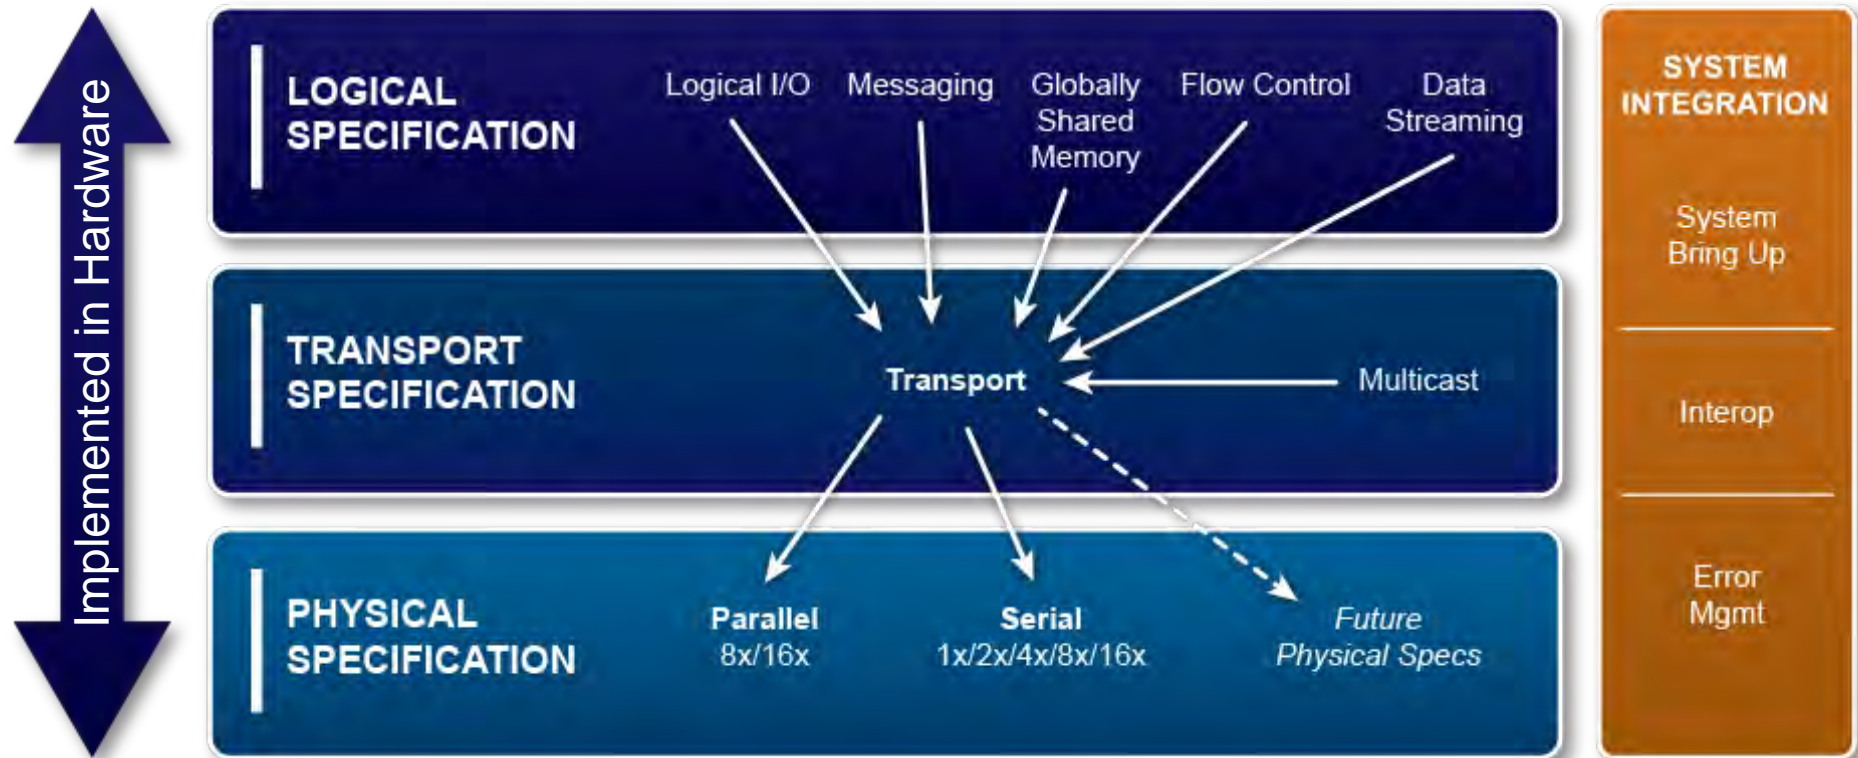
- Proven technology > 10 years of market deployment
- Supported by major CPU, DSP, NPU, FPGA & system vendors
- 10Gbps/lane Specification (10xN) released Q4 2013
- 10xN 3.1 Spec released Q3 2014 with increased fault tolerance for wireless & space applications
- 25Gbps/lane (25xN) on roadmap
- Hardware termination at PHY layer
- Lowest Latency Interconnect ~ 100 ns
- Inherently scales to 10,000's of nodes



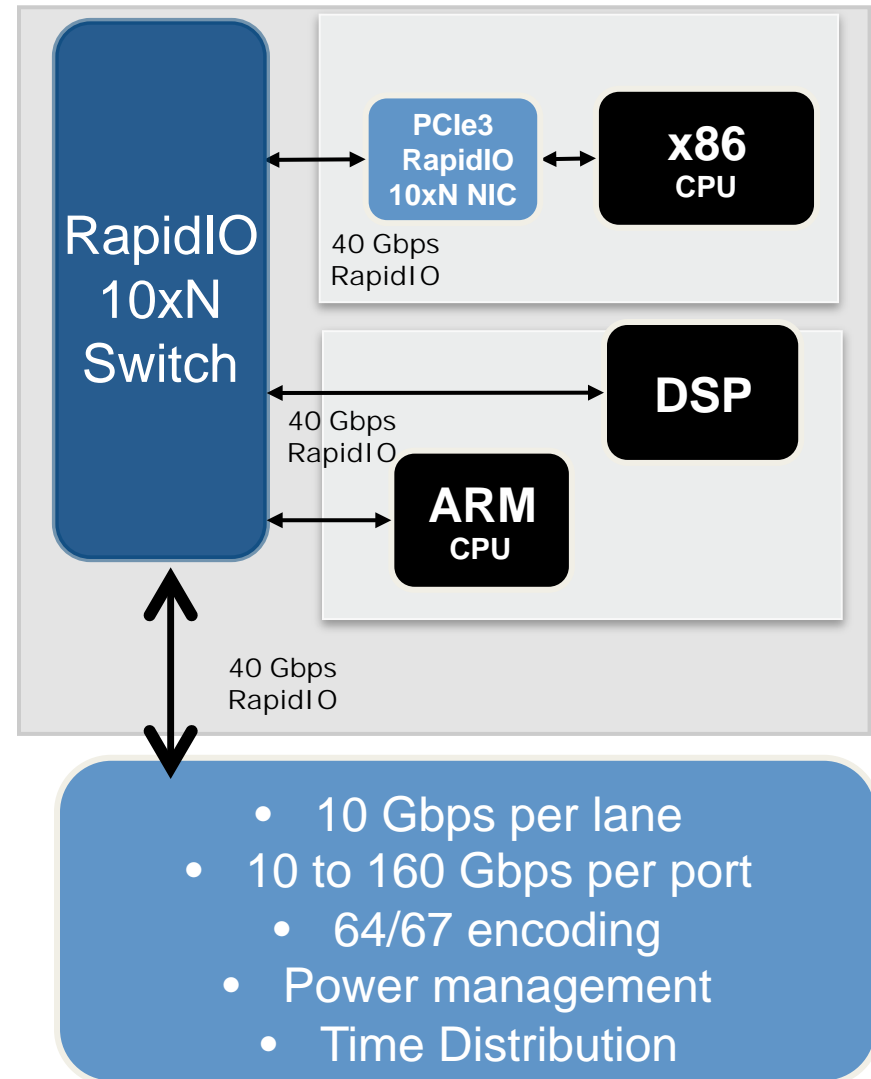
- Over 100 million 10-20 Gbps ports shipped worldwide
 - 100% 4G/LTE interconnect market share
 - 60% Global 3G & 100% China 3G interconnect market share

RapidIO Specification Hierarchy

Hardware Terminated Protocol Stack – no CPU overhead



- RapidIO 10xN (Gen3 Standard)
 - Rev 3.1 released Q3 2014
 - 3.0 released Q4 2013
 - 10-160 Gbps per port
 - 10.3125 Gbaud per serial lane with roadmap to 25G
- Long-reach support (100 cm through two connectors)
- Short Reach 20 cm 1 connector, 30 cm no connector
- Backward compatibility with RapidIO 6xN (Gen2)
- Lane widths of x1, x2, x4, x8, x16
- Speed granularity from 1.25, 2.5, 3.125, 5, 6.25, 10.3125 Gbaud






Strong and Growing Ecosystem

- Supported by major CPU & NPU vendors
- Supported by major DSP vendors
- Supported by major FPGA vendors
- Multiple IP vendors offering products
- Multiple protocol analyzer, network management tools vendors
- Over 100 million 10-20 Gbps ports deployed
- We estimate total annual RapidIO enabled semiconductor sales to be > \$1 billion (DSP, Processor, Switch, FPGA, ASIC)
- Strong market penetration in:
 - Data Center HPC & Analytics
 - Wireless Infrastructure
 - Industrial Automation
 - Mil/Aero

RapidIO Strong and Growing Ecosystem



- The ARM 64-bit Coherent Scale Out over RapidIO Task Group shall be responsible for developing a specification for multi SoC / core coherent scale out of ARM 64-bit cores with the following functionality:
 - coherent scale out of a few 10s to 100s cores & 10s of sockets
 - ARM AMBA[®] protocol mapping to RapidIO protocols
 - AMBA 4 AXI4/ACE mapping to RapidIO protocols
 - AMBA 5 CHI mapping to RapidIO protocols
 - Migration path from AXI4/ACE to CHI and future ARM protocols
 - support for GPU/DSP floating point heterogeneous systems
 - HW hooks and definition to support RDMA, MPI, secure boot, authentication, SDN, Open Flow, Open Data Plane, etc
 - Other functionality as necessary to for performance critical computing
 - support Data Center, HPC and Networking Infrastructure system development and deployment





System Challenges

Networking and Computing Infrastructure



Multi-Processor
Embedded Interconnect

Switched | Scalable | Low Latency | Reliable

ANY TOPOLOGY
ANY PROCESSOR
OPEN STANDARD



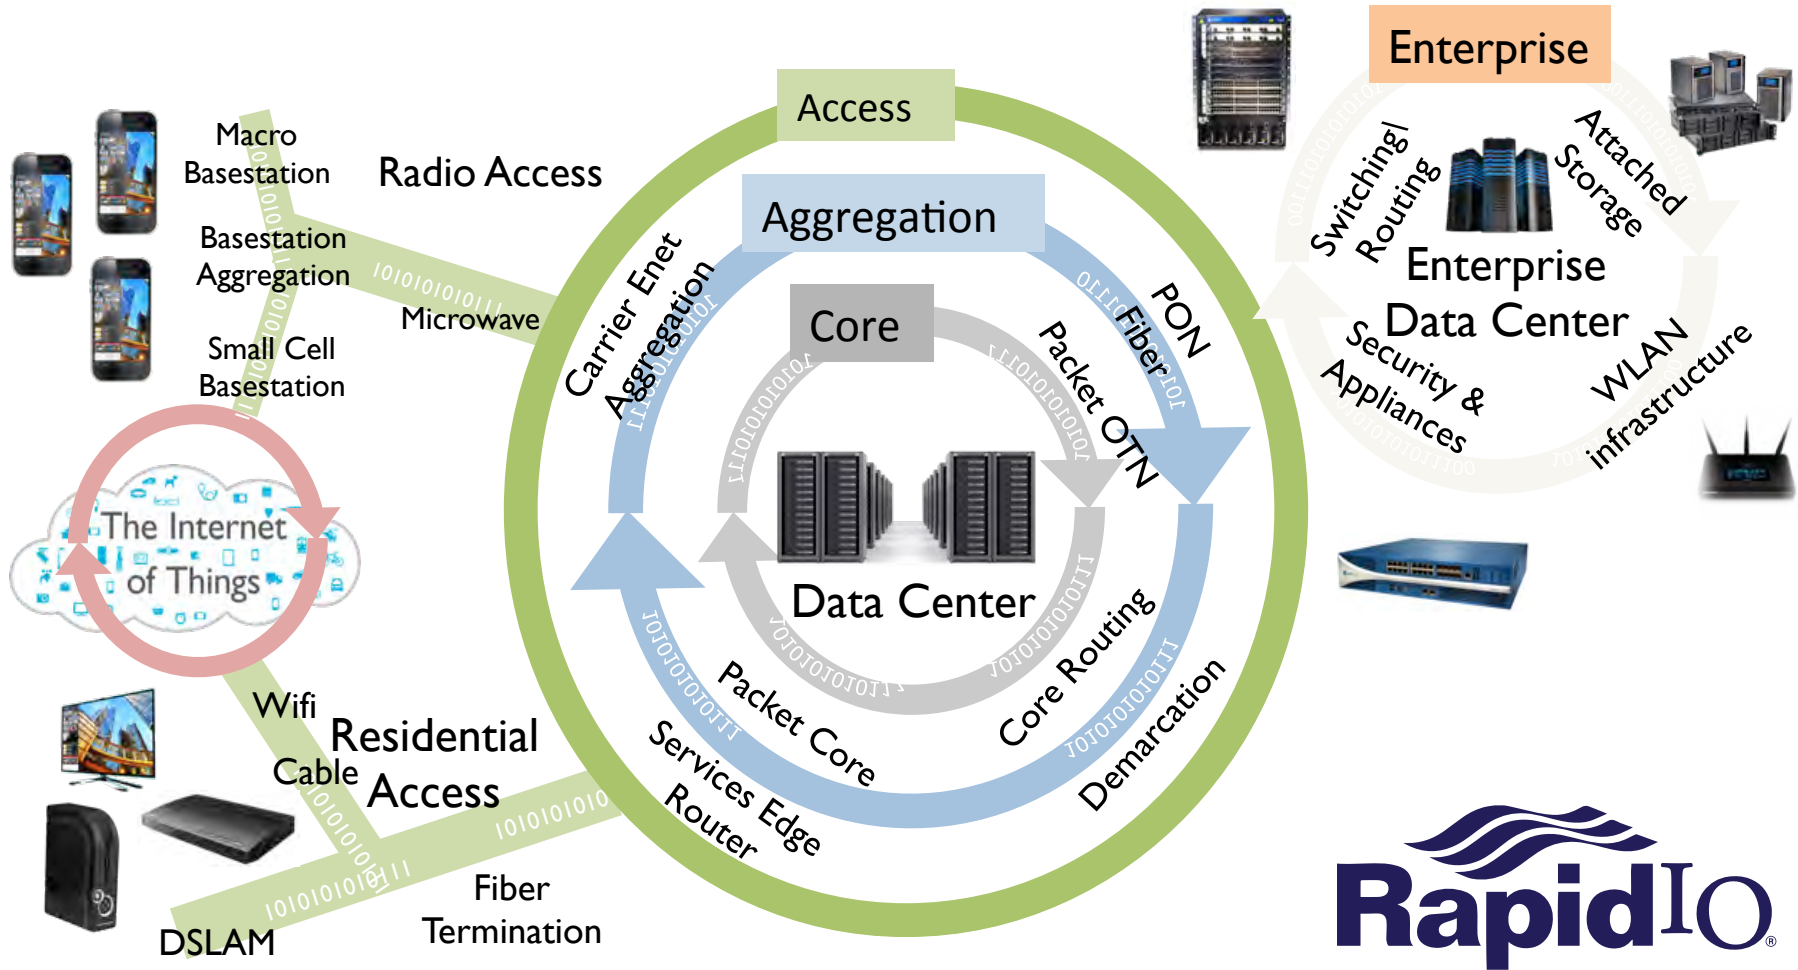
10 Gbps

20 Gbps


40 Gbps

100+ Gbps

WIRELESS INFRASTRUCTURE | SERVER | HPEC | IMAGING | AEROSPACE | INDUSTRIAL



Application specific boxes transitioning to converged, scaled heterogeneous nodes



Many power constrained form factors

- Small cell: ~5W SoC
- Macro cell: ~25W SoC

Data center power budget

- ATCA blade: ~200W

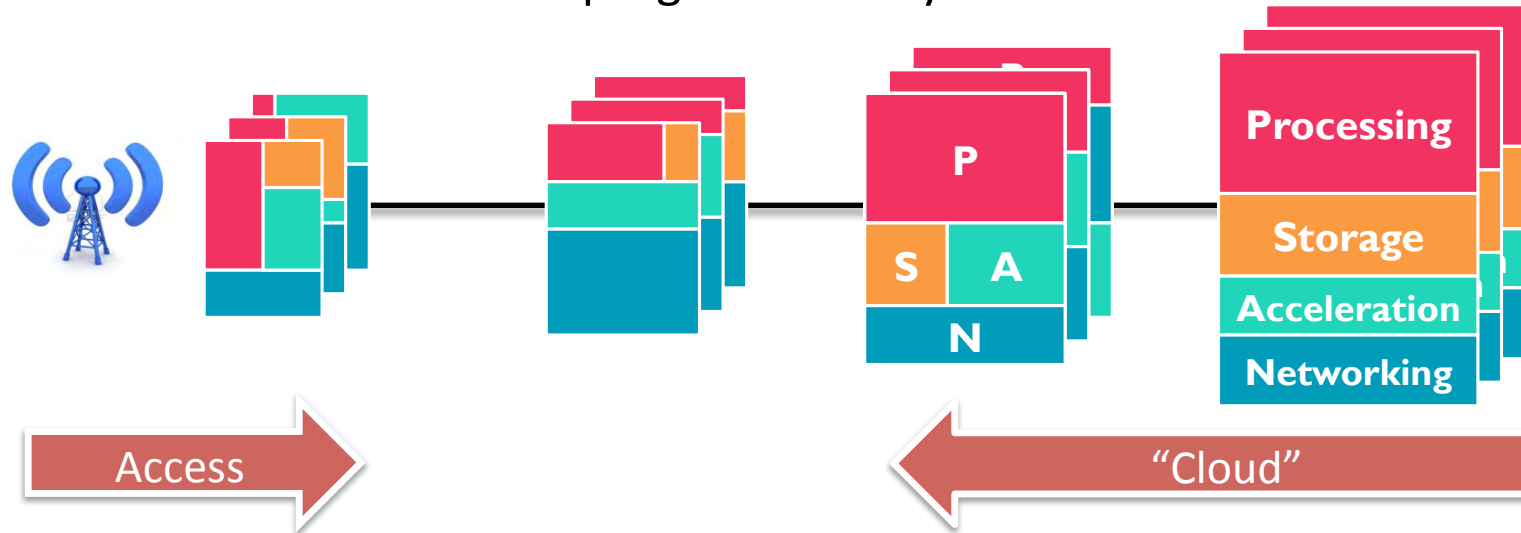
Move to general purpose processors
Virtualisation

Separation of operational plane (hardware) and control plane

Consolidation (different boxes/subsystems within a box)

Flexibility for evolving workloads

Pool of scalable heterogeneous nodes with increased visibility and programmability



Access to distributed intelligence for flexibility of workloads, traffic, services
 Heterogeneous platforms to meet diversity of physical requirements,
 workload requirements



DCCN - Data Center Compute & Networking System

Open Source Performance Critical Compute Platform



Multi-Processor
Embedded Interconnect

Switched | Scalable | Low Latency | Reliable

10 Gbps

20 Gbps

40 Gbps

100+ Gbps

ANY TOPOLOGY
ANY PROCESSOR
OPEN STANDARD




WIRELESS INFRASTRUCTURE | SERVER | HPEC | IMAGING | AEROSPACE | INDUSTRIAL

- **RapidIO DCCN System** - Open Industry wide collaboration with several semiconductor, software and board/systems vendors participating
- **RapidIO** is a best-in-class open community, low-latency lossless fabric with over 100M 10-20 Gb/s ports deployed with 5Nines reliability
- DCCN System is a processor agnostic, heterogeneous compute platform with:
 - 300 Gb/s **RapidIO** Switching
 - 268.8 GFlops/1U - example configuration using 4-core Intel i7-3612QE 2.1 GHz
 - 64 GB of DDR3-1600 memory




Visit RapidIO.org to learn more

DCCN System in 19" Rack



- Texas Instruments Keystone II ARM + DSP with native RapidIO
- 1x TI 66AK2H12: 4 ARM Cortex-A15 and 8 C66x DSP
- 2x TI C6678: 16 C66x DSP
- 518 GLOPS at < 50W typical
- Freescale P4080 and P5020 with native RapidIO
- Octal e500 or dual 64-bit e5500 Power Architecture
- Up to 16GB DDR3 memory
- 2x RapidIO quad lane to AMC




- Intel i7 quad core with RapidIO to PCIe Bridging
- 16-32 Gbps backplane connectivity with IDT Tsi721
- PCIe to RapidIO NIC only 2W, 300 ns latency
- Native RapidIO 2x TI C6678 Processors
- 640 GMAC + 320 GFLOPS (2 CPUs 8 cores fixed and floating point)
- 40 Gbps backplane connectivity
- No NIC zero processor latency

CONCURRENT TECHNOLOGIES 



Comm (())
Agility




ARM + DSP HPC over RapidIO

- HPC System based on TI Keystone I & II multicore processors
- 576 ARM Cortex A15 MPCores & 3456 C66x CorePac DSPs
- 66.3 TFLOPS single precision & 18.6 TFLOPS double precision
- 14.8 Tb/s aggregate interconnect throughput
- OpenMPI over RapidIO
- OpenMP / OpenCL programming model
- Linux OS

Intel i7 Hadoop Cluster on RapidIO DCCN

- 4 x Concurrent Technologies AMC cards featuring dual quad core Intel® Core™ i7 processors providing aggregate of 32 Cores per server
- Concurrent Technologies highly optimized RapidIO Software stack
- Open Standards based Data Center Compute and Networking (DCCN) platform
- 20 Gbps/port Low Latency High Throughput RapidIO Fabric Apache™ Hadoop® enabling Big Data processing





VITA 78 SpaceVPX & NGSIS

Next Generation Space Interconnect Standard



Multi-Processor
Embedded Interconnect

Switched | Scalable | Low Latency | Reliable

10 Gbps

20 Gbps

40 Gbps

100+ Gbps

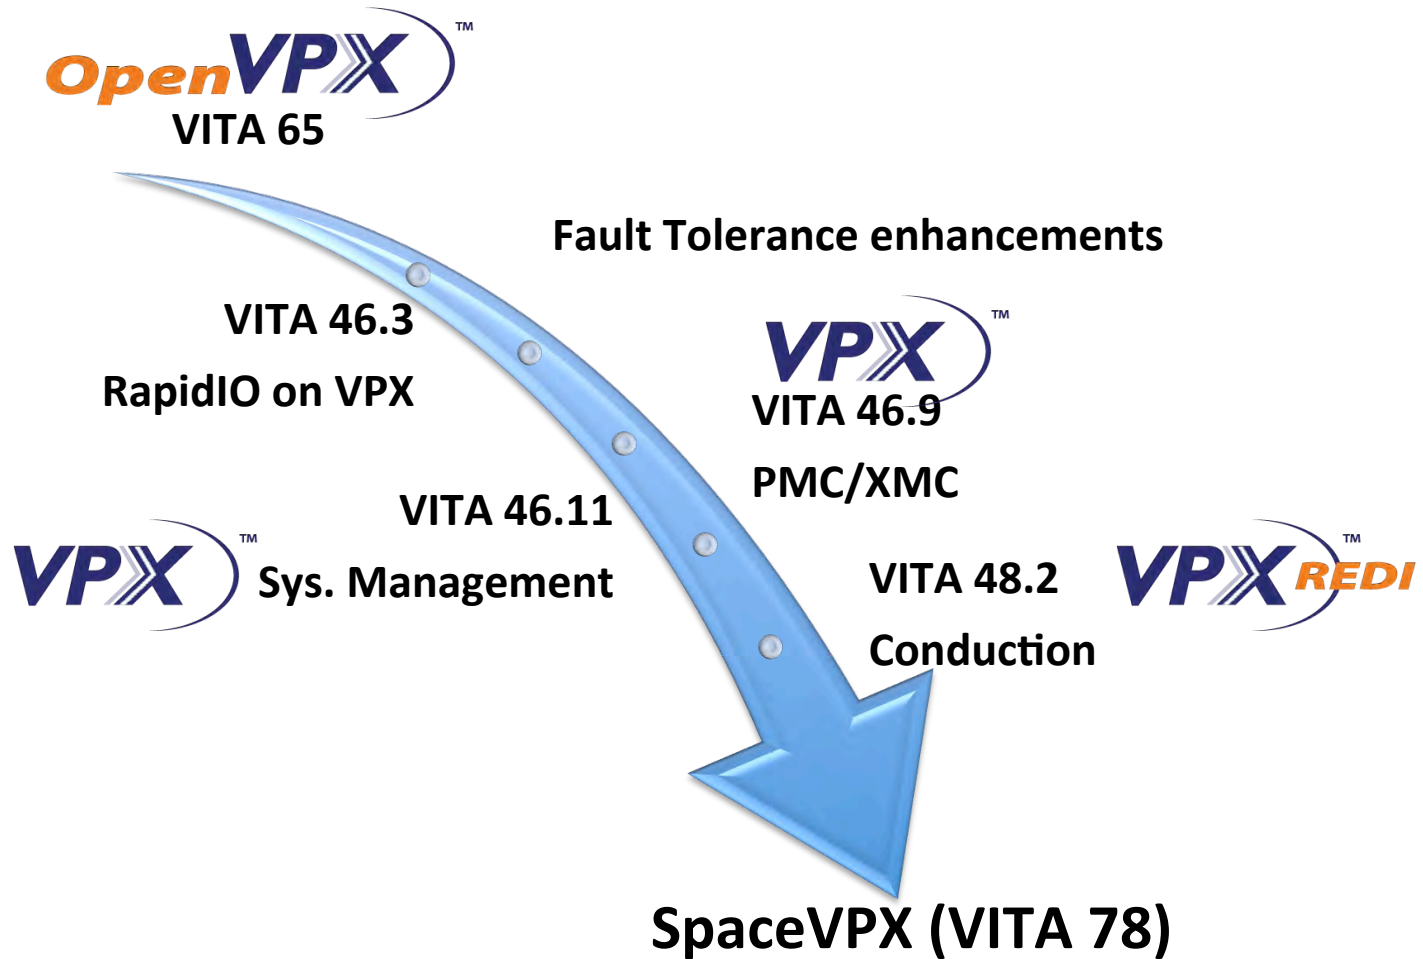
ANY TOPOLOGY
ANY PROCESSOR
OPEN STANDARD



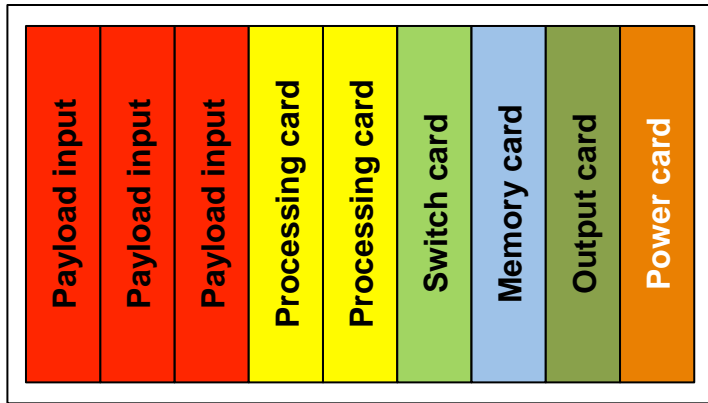
WIRELESS INFRASTRUCTURE | SERVER | HPEC | IMAGING | AEROSPACE | INDUSTRIAL

- In 2012, the NGSIS team selected the RapidIO protocol for adoption in NGSIS systems
- Working with RapidIO.org, a “Part S” Task group was formed to develop a unique features for use in space applications
- Focus was on adding additional robustness and fault tolerant features to the RapidIO specification
- Because many of these features apply to other applications as well, the Part S features have been built into the existing RapidIO 10xN specification as Revision 3.1


- Honeywell
- BAE Systems
- Harris
- Boeing
- Lockheed Martin
- Northrop Grumman
- SEAKR Engineering
- L3 Communications
- ELMA Bustronics
- Aerospace Corporation
- TE Connectivity
- Raytheon
- Smiths Connectors
- Amphenol
- Curtiss-Wright
- NASA
- NASA/JPL
- NRL
- SMC-XR
- Microsemi
- Aeroflex
- Freescale
- Xilinx
- IDT
- Texas Instruments
- Mobiveil
- Orbital Sciences
- IEH



- Protocols: RapidIO, SpaceWire, I2C, etc.

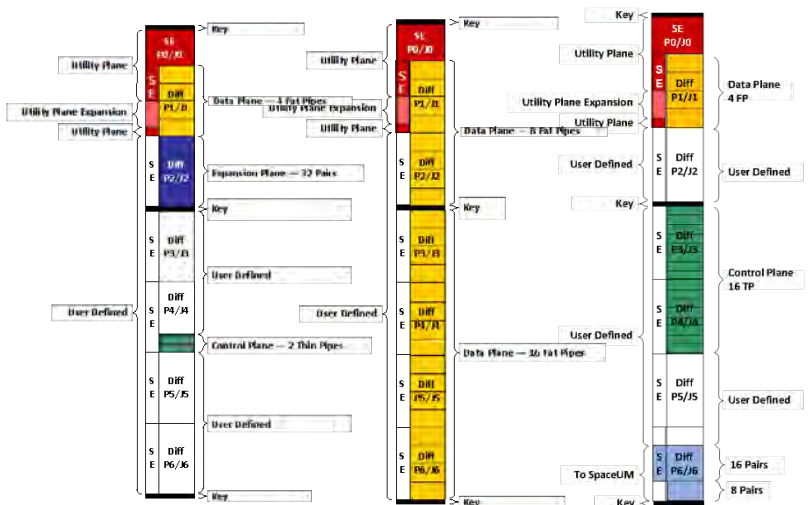


Payload

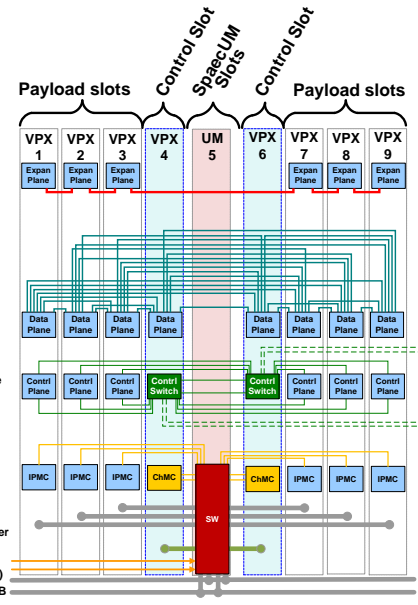


Bus

MILSATCOM System



Slot numbers are logical, physical slot numbers may be different



RapidIO 10xN 3.1 Enhancements

- Structurally Asymmetric links
- Degradation of 4x/2x/1x port to 2x operation on lanes 2 and 3.
- Require all of Part 8 Error Management Extensions
- Multicast Extensions spec for dev8 and dev16
- Time to Live functionality in switches
- MECS Time Synchronization
- Support for software assisted error recovery
- Performance Diagnostics (PRBS generation)



RapidIO based HPC Systems

HP Moonshot / IDT HPAC Lab / OCP HPC & RapidIO



Multi-Processor
Embedded Interconnect

Switched | Scalable | Low Latency | Reliable

ANY TOPOLOGY
ANY PROCESSOR
OPEN STANDARD



10 Gbps > 20 Gbps > 40 Gbps > 100+ Gbps

WIRELESS INFRASTRUCTURE | SERVER | HPEC | IMAGING | AEROSPACE | INDUSTRIAL



- 2D Torus RapidIO unified fabric
- up to 45 m800 cartridges capable of providing 5Gbps per lane connections in each direction to its north, south, east and west neighbors
- highest density DSP solution in an industry standard infrastructure in the market today:
 - 1,440 C66x DSP cores
 - 760 ARM A15 cores
 - up to 11.5TB of storage in a single Moonshot chassis
 - all connected via a 5Gbps per lane RapidIO unified fabric

Order out of chaos



“The HP Moonshot ProLiant m800’s combination of ARM and multi-core DSPs with high-speed, low-latency networking and tiered memory management creates a very energy efficient, extremely capable parallel processing platform with a familiar Linux[®] interface. It’s a truly new approach to bringing scale-out design ‘inside the box,’ and breaks barriers between HPC and Enterprise technology.”

Ryan Quick, Principal Architect,
PayPal

High Performance Computing and Data Analytics with Low Latency

Objectives

RapidIO based Platform for Applied Architecture and Product Development

Collaboration on Existing and Future IDT Products

Public/Private Cloud Platform Enables Collaboration

Best-in-class \$/Watt/MHz System



Architecture/Technology

Heterogeneous Computing

Scale-Out System

Fault-tolerant Transport

100 nsec interconnect latency

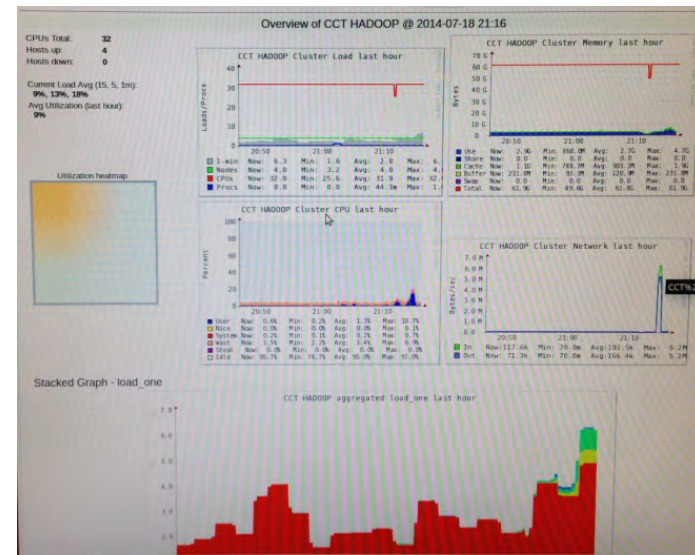


High-Performance Computing and Analytics Lab

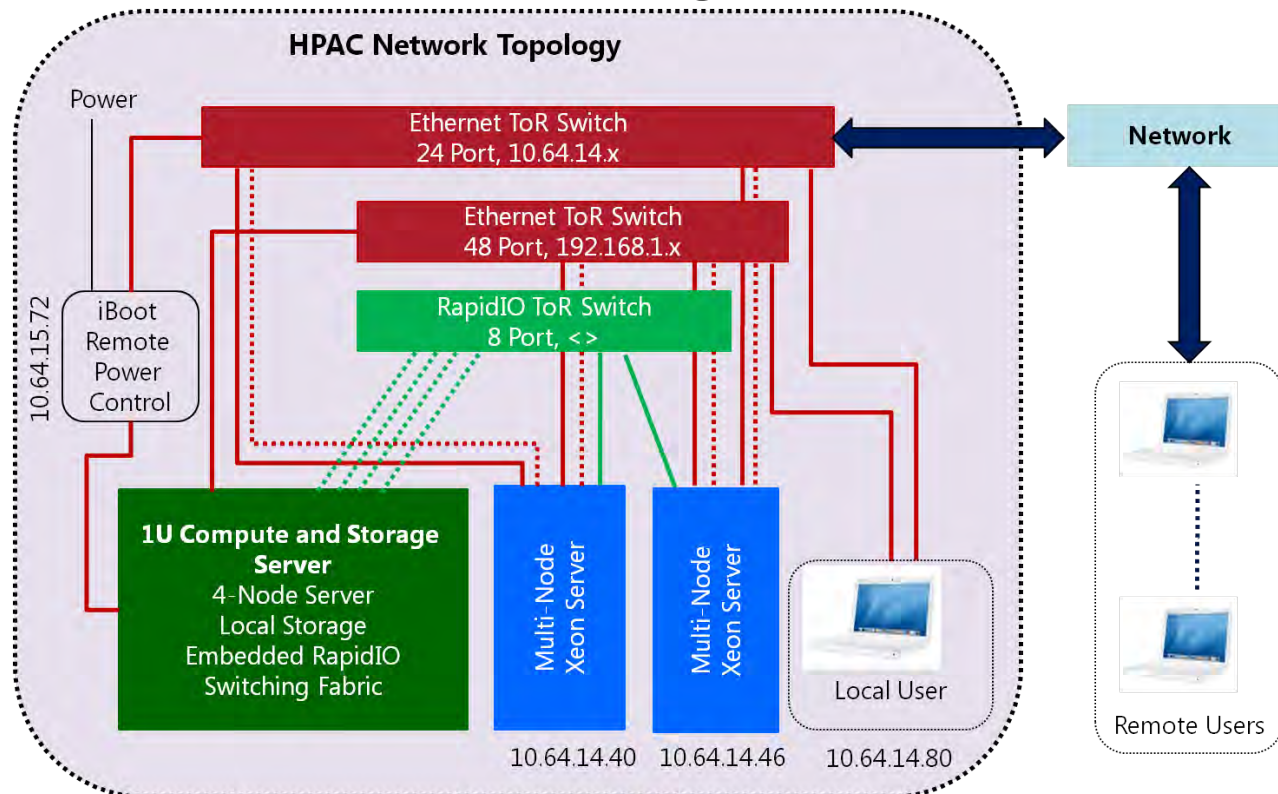
- Top Green 500 caliber energy efficiency
- 100 ns latency
- Cluster scales to thousands of nodes
- Mission-critical reliable network
- Industry-wide collaboration
- Open platform

IDT RapidIO® Connected

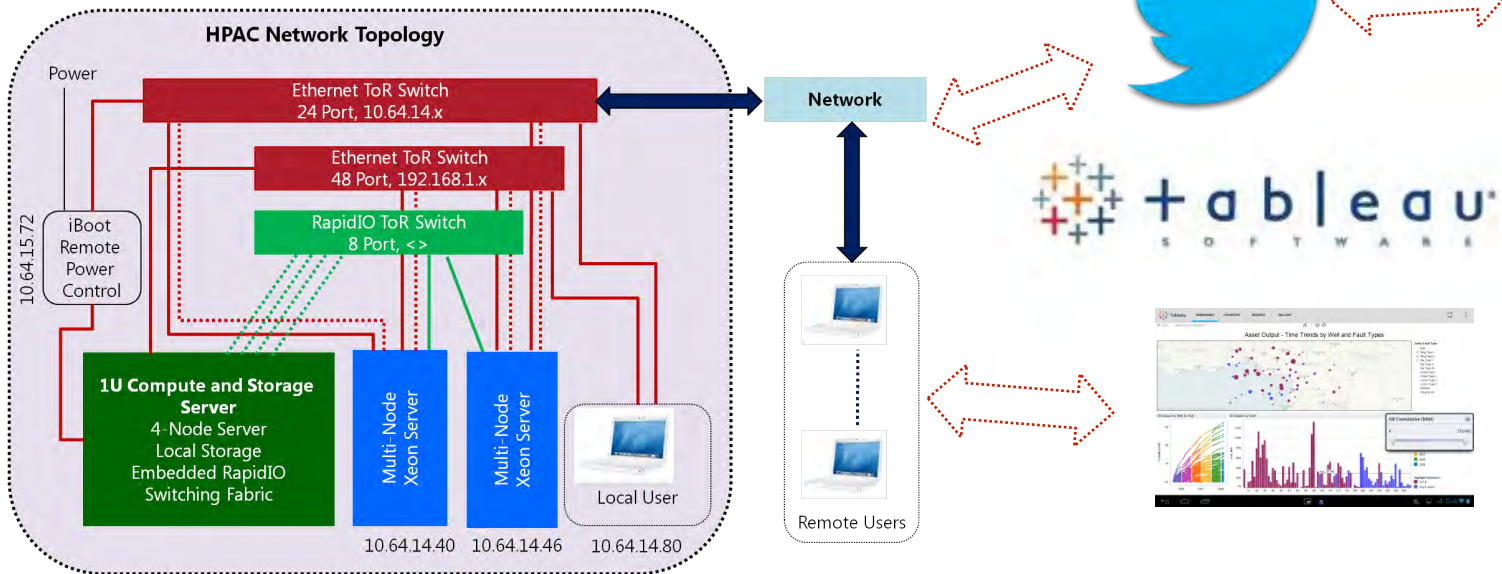
- Analytics Hardware and Software
 - 300 Gbps in-chassis RapidIO unified fabric
 - 92.6% Bandwidth Utilization using TCP/IP based high performance data-path SW



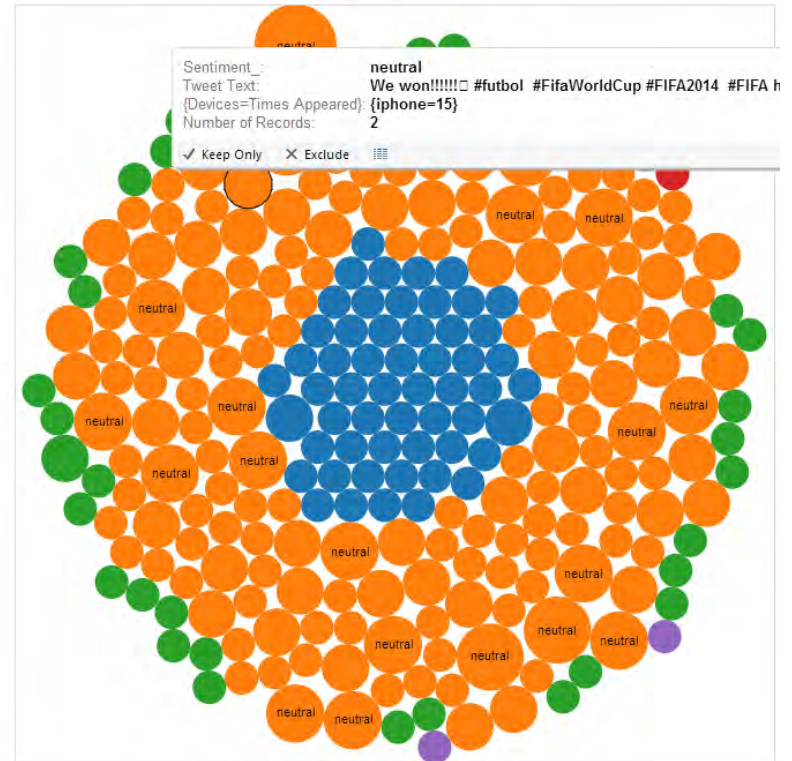
- IDT Partners / Customers / Developers Play Ground...
 - Computing, Storage, Networking as a Service for Technology Development
 - Private / Public Cloud Paradigm



- Analyze User sentiment of 2014 World Cup Final (Argentina vs Germany)
 - Users impression of the game over Wireless Access Network
 - Collect Twitter Data over the network
 - Perform Data Analytics using Hadoop cluster
 - Present User Impressions using Tableau



- World Cup Final 2014 Game Analytics
 - Plotted using Tableau Public Software
 - Identify Number of users and Smart Phone/Source Type
 - Identify Set of Users' feelings {Positive, Neutral, Very Positive, Negative, Very Negative}



- Low latency analytics critical to HPC and Supercomputing and other computing applications as well (wireless base stations)
- OCP invited IDT to form an OCP HPC group with a mandate to create Open, latency sensitive, energy efficient board level & rack level platforms
- Multi phase project at OCP with a path to Open industry interconnect silicon
- Several industry players will table open spec based on RapidIO 10xN
- RapidIO DCCN platform to be submitted for OCP certification





Multi-Processor
Embedded Interconnect

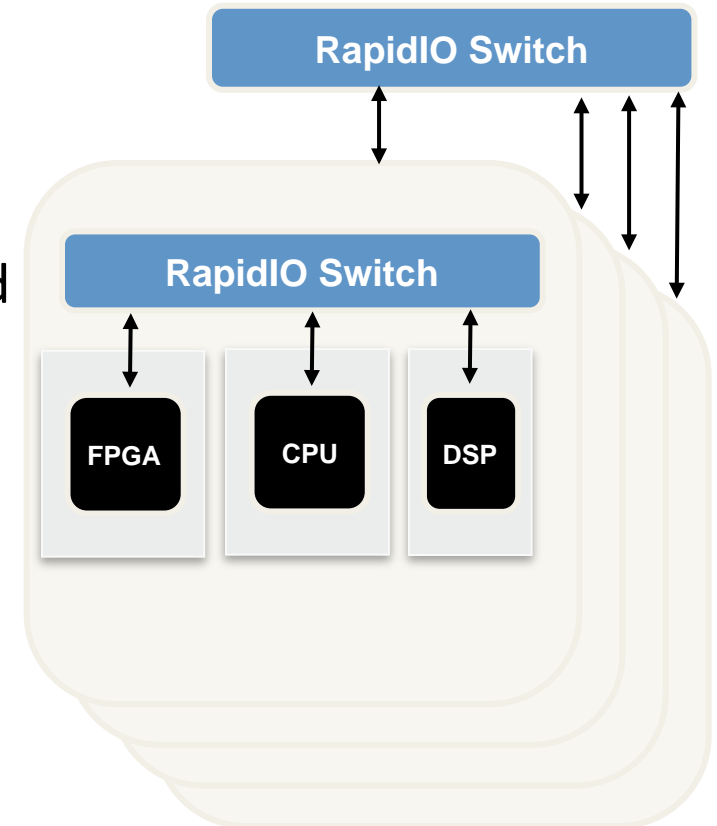
Switched | Scalable | Low Latency | Reliable

ANY TOPOLOGY
ANY PROCESSOR
OPEN STANDARD

10 Gbps > 20 Gbps > 40 Gbps > 100+ Gbps

WIRELESS INFRASTRUCTURE | SERVER | HPEC | IMAGING | AEROSPACE | INDUSTRIAL

- Proven technology > 10 years of market deployment
- Supported by major CPU, DSP, NPU, FPGA & system vendors
- 10Gbps/lane Specification (10xN) released Q4 2013
- 10xN 3.1 Spec released Q3 2014 with increased fault tolerance for wireless & space applications
- 25Gbps/lane (25xN) on roadmap
- Hardware termination at PHY layer
- Lowest Latency Interconnect ~ 100 ns
- Inherently scales to 10,000's of nodes



- Over 100 million 10-20 Gbps ports shipped worldwide
 - 100% 4G/LTE interconnect market share
 - 60% Global 3G & 100% China 3G interconnect market share