Integrated Device Technology

RapidIO 50-100 Gbps Technologies for MicroTCA Systems

December 2015

Presenter: ruben.aszkenasy@idt.com

What and Where is this?

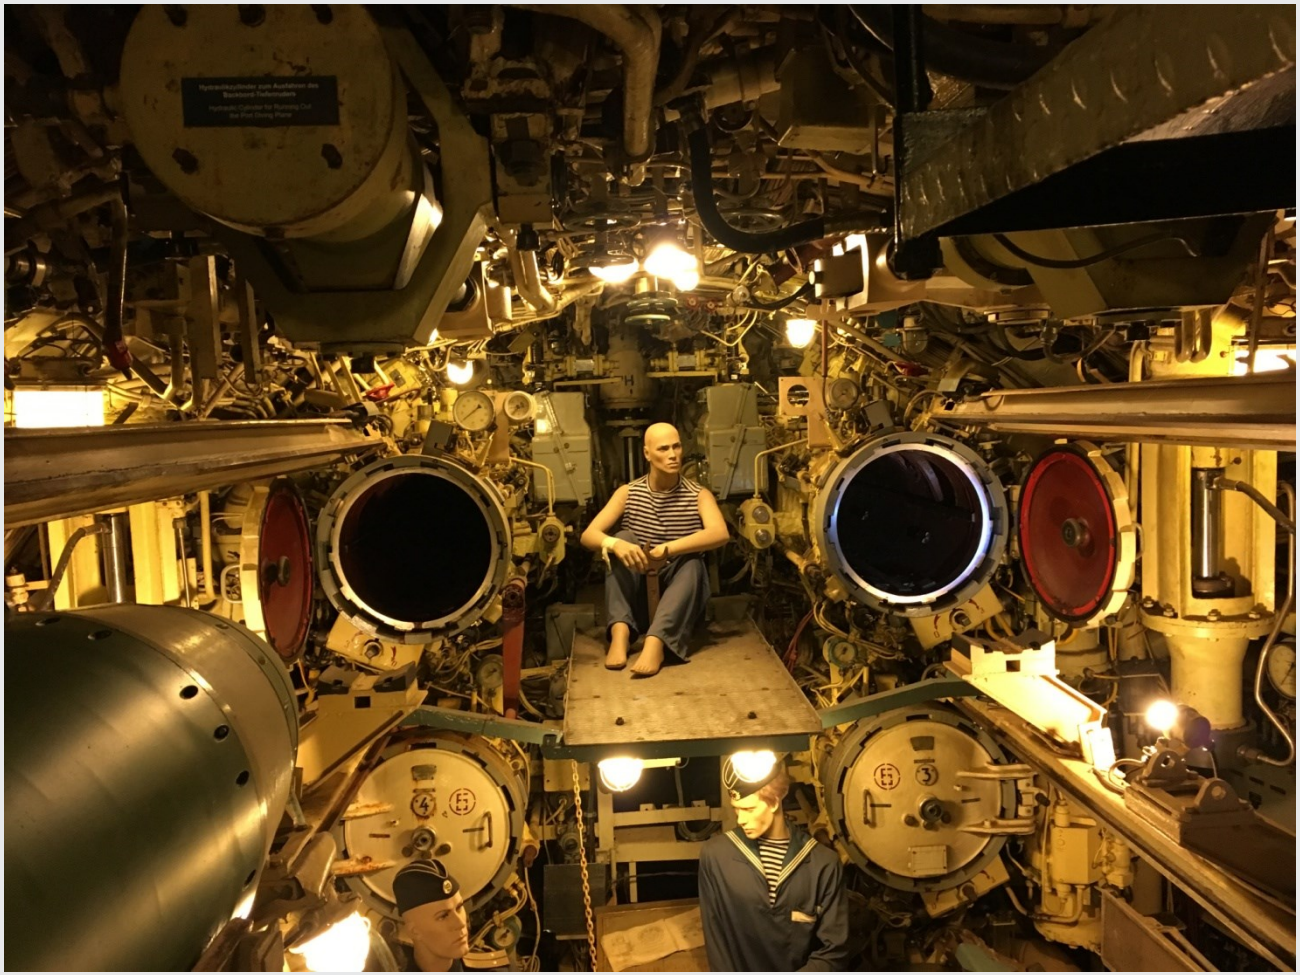

RapidIO 50-100 Gbps Technologies for MicroTCA Systems

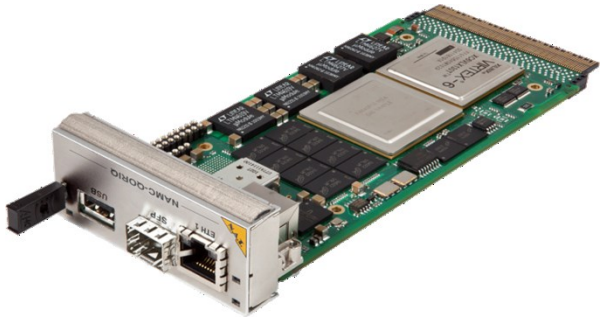

- RapidIO Ecosystem
- Protocol Attributes
- RapidIO 10xN Features
- Systems Architectures
- AMC ecosystem and solutions for Heterogeneous Computing
- Systems Solutions at CERN and Orange Telecom

IDT Company Overview

Founded	1980
Workforce	Approximately 1,500 employees
Headquarters	San Jose, California

#1 Serial Switching – 100% 4G Infrastructure with RapidIO

#1 Memory Interface – Industry Leader DDR4

#1 Silicon Timing Devices – Broadest Portfolio

800+ Issued and Pending Patents Worldwide

Mixed-signal application-specific solutions

- **20/40/50Gbps per port – 6.25/10/12.5 Gbps lane**
- Embedded RapidIO NIC on processors, DSPs, FPGA and ASICs.
- Hardware termination at PHY layer: 3 layer protocol
- Lowest Latency Interconnect ~ 100 ns
- Inherently scales to large system with 1000's of nodes

- **Over 14.5 million RapidIO switches shipped**
- **> 2xEthernet (10GbE)** Over 85 million 10-20 Gbps ports shipped
- **100% 4G interconnect market share**
- 60% 3G, 100% China 3G market share

Live Data from Wireless to the Cloud

The Real-Time Usage Rising

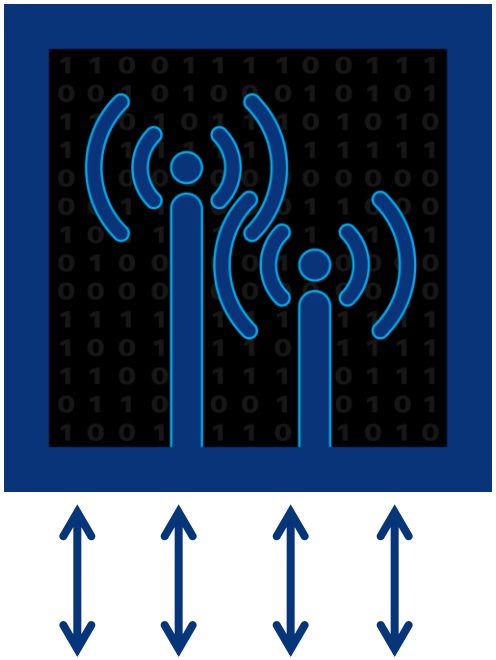

Networks and Data Centers must THINK FAST

- Millions of users
- Billions of devices
- More data per user
- Unstructured Data
- Video/Text/Image

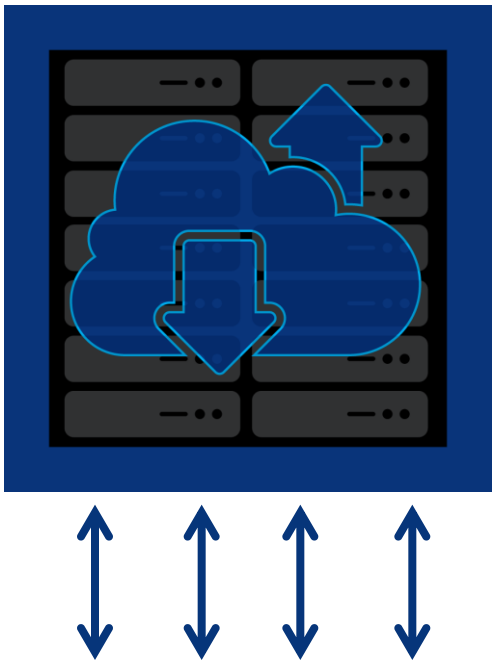

Social Media, Finance, Health, Video, Audio, Cloud Compute, Transportation, Military C3I

RapidIO Markets Today

RapidIO

- Lowest latency interconnect
- High scalability
- Industry standard fabrics
- Energy efficiency
- Reduced TCO

Over **400** production customers across markets

Wireless

Defense and Aerospace

Real-time Image Analytics

Industrial Automation

Data Center

Big Data Analytics

Scalable Storage

HPC

Industrial, Military & Aerospace

- RapidIO Fabrics deliver key performance metrics such as:
  - best in class switch cut through latency of ~100ns
  - lossless data flow
  - small footprint and lowest overall power consumption
  - High reliability for mission critical markets
- RapidIO.org Community works closely with the PICMG for ATCA/MicroTCA/AMC standardization

RapidIO Fabrics deliver highly reliable, fault tolerant performance to demanding Mil/Aero applications

Supported by Large Eco-system

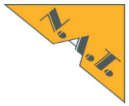

RapidIO 10xN in Wireless and C-RAN

3G → 4G → 5G

- Today 100% of the 4G OEMs use RapidIO for baseband interconnect
- 4G technologies are driving need for inter processor communication
- Wireless and Server worlds are converging
- IDT is developing S-RIO 10xN Switches, bridges and endpoint IP @ 50+ Gbps per port
- Peer to Peer Scalable Interconnect for Wireless, Cloud, Imaging, Military, Industrial and Medical
- Scales to 100 Gbps per port with SerDes Upgrade

- 10.3125 Gbaud per lane, 50 Gbps per port embedded interconnect
- 100 ns Latency with 5x effective bandwidth of 10 GigE for embedded systems
- <300mW per 10 Gbps of data

RapidIO Specification Mapped to Devices

Logical

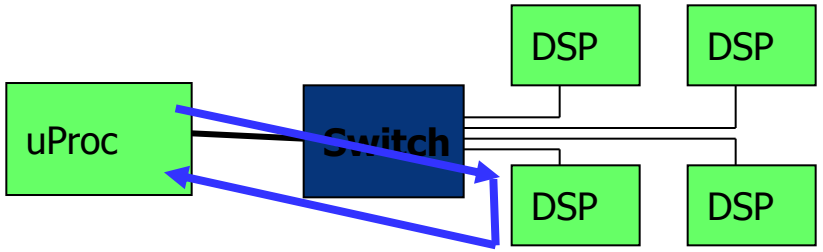

Transport

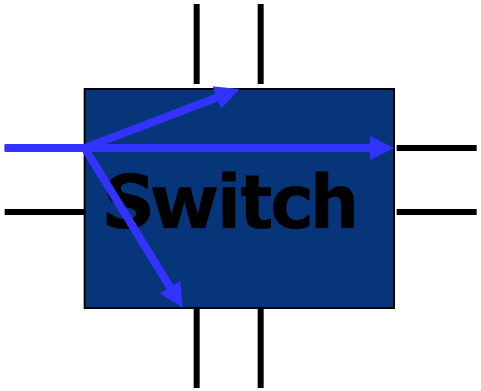

Physical

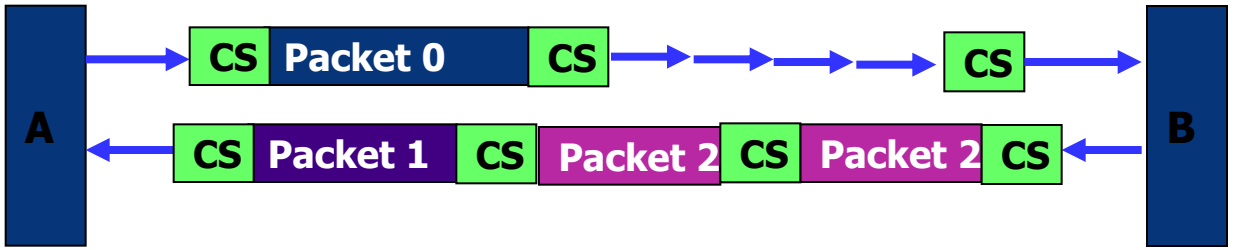

RapidIO for Fault Tolerant Systems

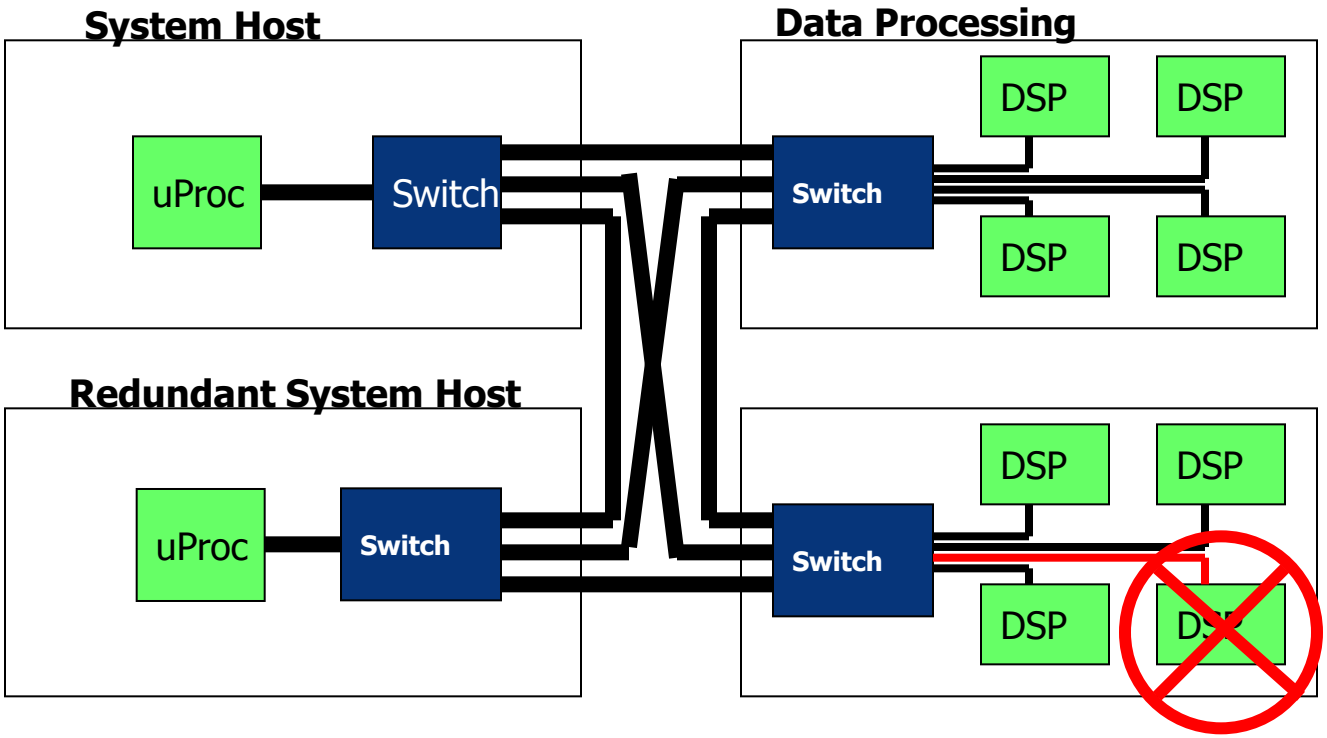

- Packets destined for the faulty endpoint must be discarded
  - Refer to Part 8 Error Management Extensions Section 1.2.4, table 1-1
  - Timeouts can also be used to discard packets
- Hardware must detect and rapidly notify the System Host of a failure
  - The hardware must begin discarding packets to prevent system congestion/failure
- The System Host must receive notification, then diagnose, isolate and allow replacement of the faulty hardware

Embedded Interconnect Needs

Interconnect Requirements	RapidIO	Details
Low Latency		switch silicon: ~100 nSec memory to memory : < 1 uSec
Scalability		support any topology, 1000's of nodes, true peer-to-peer
Integrated HW Termination		integrated into SoCs and guaranteed, in-order delivery without software overhead
Power Efficient		3 layers terminated in hardware, Integrated into SoC's
Fault Tolerant		supports hot swap and fault tolerance
Deterministic		guaranteed, in order delivery with deterministic flow control
Top Line Bandwidth		10-25 Gbps per lane

Why RapidIO 10xN in Embedded

Bandwidth and Per Port Economics			
System Requirement	**Embedded 10G Ethernet**	**RapidIO Gen 2**	**RapidIO 10xN**
Switch per-port performance raw data rate	**10 Gbps**	**20 Gbps**	**50 Gbps**
End to end packet termination	**>10 ms**	**~1-2 us**	**~1-2 us***
Messaging performance	**Moderate**	**Excellent**	**Excellent**
Volume pricing $ per 10Gbps	**$ 8**	**< $4**	**< $3**
Typical Power Per 10 Gbps	**>1W**	**<0.4W**	**<0.3W**

- * The S-RIO latency numbers allow for a bridge (or external NIC). For applications using embedded NICs, we would see sub micro-second latencies.
- S-RIO 10xN = 50 Gbps native fabric interface, without the need for external NIC
- 40 Gbps native fabric interconnect with S-RIO 10xN, ~ 5x 10GbE actual throughput
- **Best cut through latency** 100 ns, 10 GigE cut through latency of multiple hundreds of ns
- **Guaranteed delivery** mechanisms

Clustering Fabric Needs

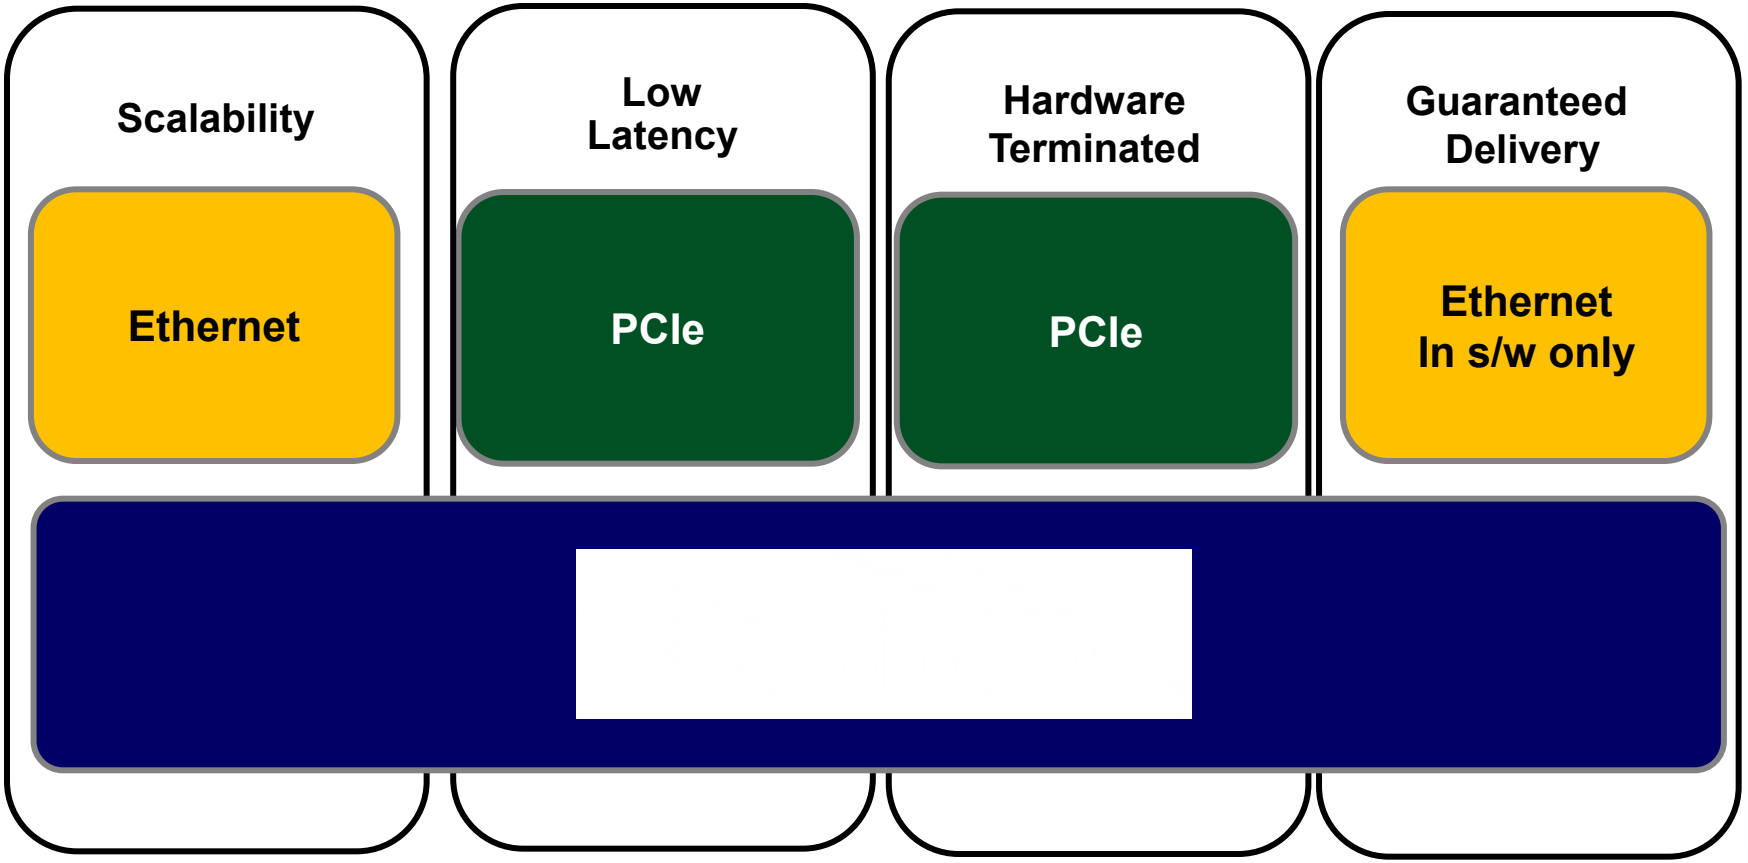

RapidIO Interconnect combines the best attributes of PCIe and Ethernet in a multi-processor fabric

Performance Needs Low Latency Interconnect

- Chip to Chip
- Board to Board across backplanes
- Chassis to Chassis
- Over Cable
- Top of Rack
- Heterogeneous Computing

Rack Scale Fabric
For any to any compute

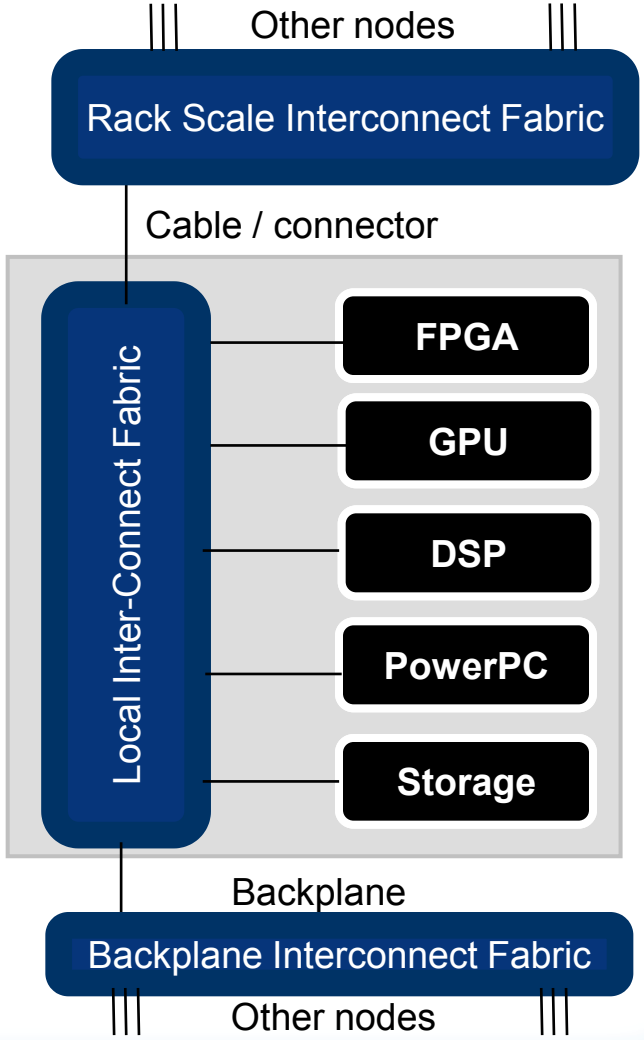

RapidIO Architectures for Heterogeneous Computing

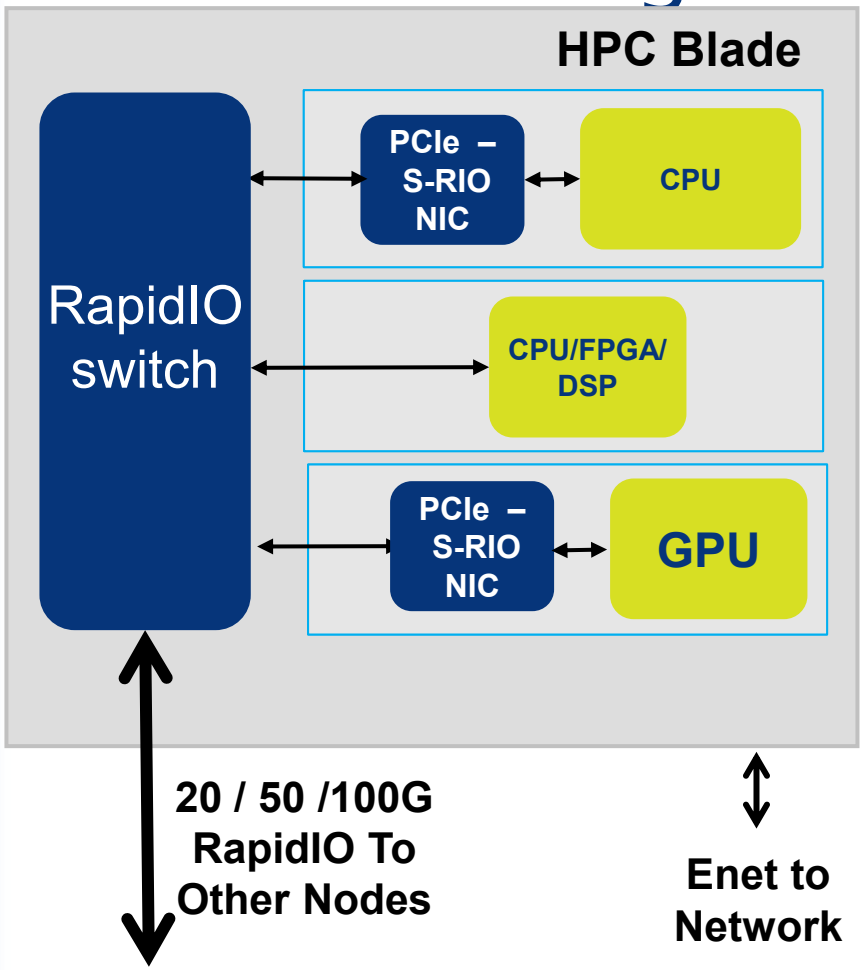

- Heterogeneous support
- With and without native RapidIO
- With GPU accelerators
- Chip-to-Chip, Board-to-Board, Chassis-to-Chassis, Over Cables, Top of Rack

High Density Computing For Embedded Apps all supported in MicroTCA

latency → scalability → Low power

AMC Ecosystem

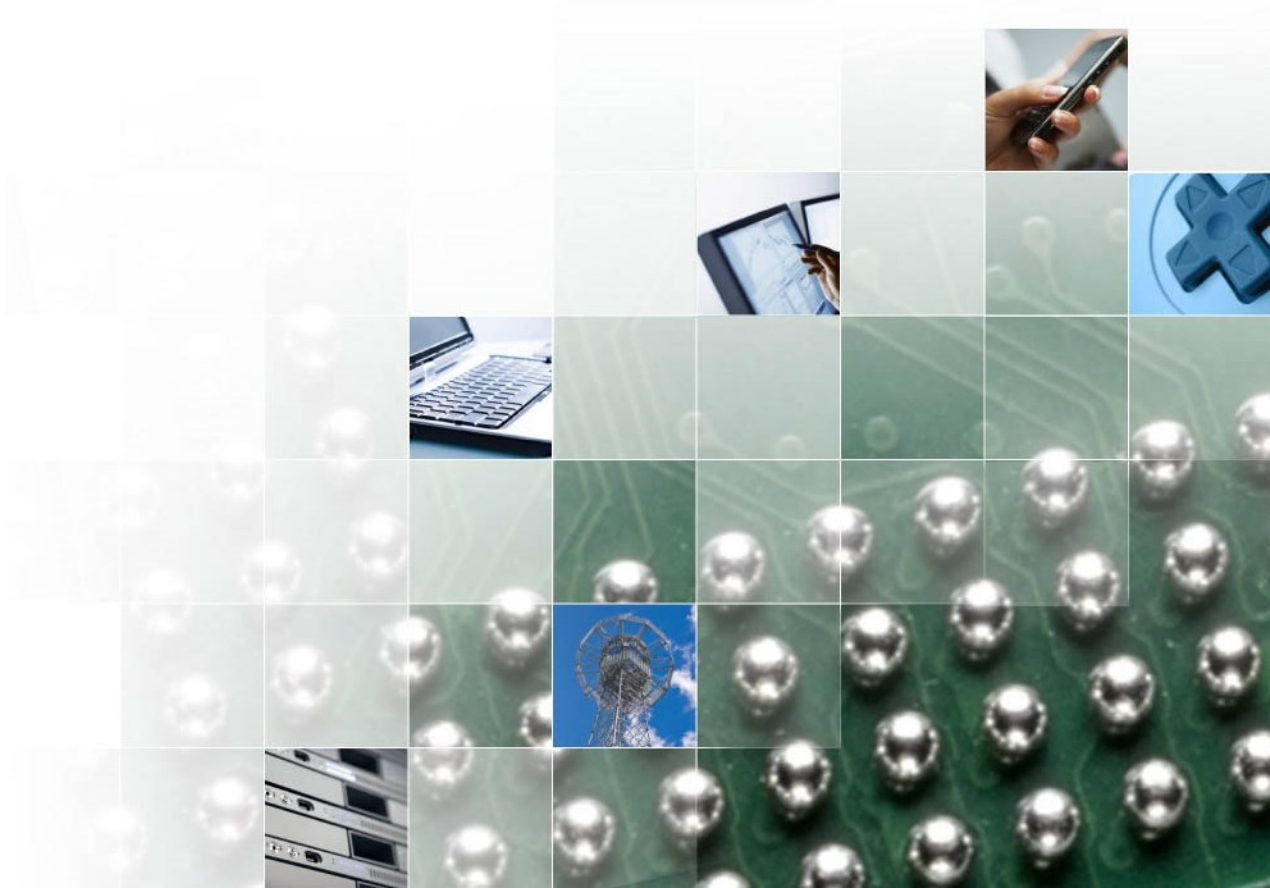

Sample RapidIO 1.3 AMCs with 10 Gbps per port

Commagility AMC
Tsi578, XilinxV4, TCI6455

Commagility AMC
Tsi578, XilinxV4, TCI6487

Mercury MTI 203 with
Tsi578, XilinxV4, TCI6482

Tsi 574 and TI 6484 (Curie)

Tsi 578 + FSL8144 DSP

Tsi 578 MSM
(MicroTCA Switching)

CPS10Q + TcCI6482

RapidIO2 in xTCA

Ecosystem for RapidIO and MicroTCA/ATCA

- Baseband Cards with DSPs with IDT switches
- Switch Cards (MicroTCA Carrier Hub)
- Antenna Interface Card (ERCA)
- PowerPC AMCs

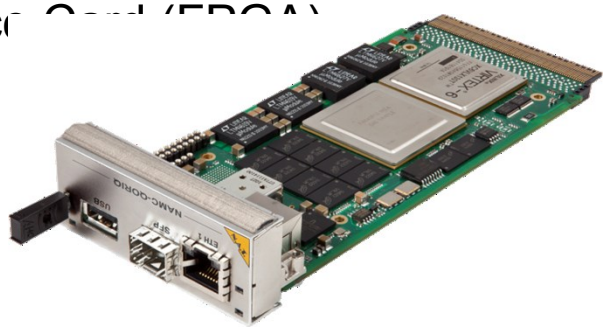

Heterogeneous Computing (DSP/ FPGA/PPC…)

Heterogeneous Computing

Proven Interoperable hardware and Industry support (Phase 1)

- Leverage industry standard AMC form-factor
- Optimize computation based on the workload

Prodrive AMC TI ARM + DSP

- ARM + DSP Processing
- Superior performance/watt Computing

NAT AMC Freescale QorIQ P4080 + Xilinx V6

- FPGA + DSP Acceleration
- High performance storage/compute

Concurrent AMC Intel CPU

- x86 Computation and Eco-system
- High Performance Analytics and Storage Solution

N.A.T. RapidIO 2 MicroTCA Carrier Hub

Provides Switching for MicroTCA System with 12x20 Gbps S-RIO

- 12x4 RapidIO ports @ 20 Gbps
- Enables customers to develop entire signal processing systems with RapidIO in the backplane
- 2x IDT CPS-1848 S-RIO Switch
- Front Panel Connectivity for chassis to chassis system level expansion
- Backward compatible with numerous S-RIO 1.3 AMCs
- AMC.4 connector compliant

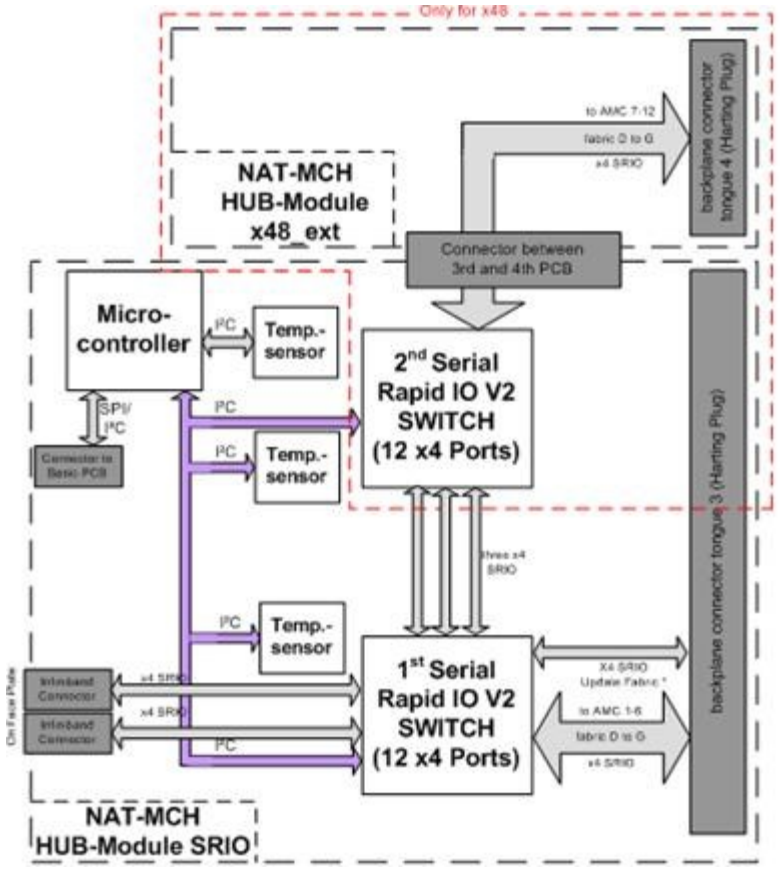

Heterogeneous Computing: GPU Compute Node

- Compute
  - 4 x Tegra K1 GPU + ARM32
- Fabric
  - RapidIO Switching
  - 4x RapidIO Gen2 NIC
  - Ethernet
- Performance
  - 1.28 Tflops/Compute (4 GPU)
  - 140 Gbps Switching Fabric
  - ~100nsec RapidIO switching latency

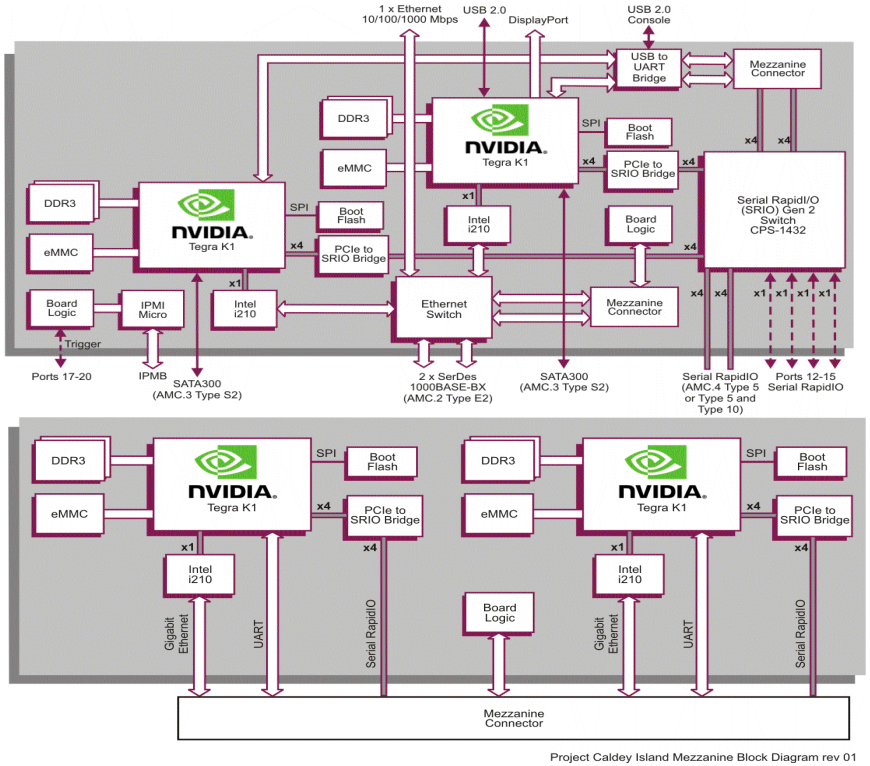

Project Caldey Island Mezzanine Block Diagram rev 01

Low Latency GPU Analytics Module
http://www.gocct.com/sheets/AG/photograph/hires/aga1xm1d.jpg
http://www.gocct.com/sheets/AG/aga1xm1d.htm

Commagility RapidIO Gen2 AMC

- Complete MAC-PHY Signal Processing Solution for wireless
  - Xilinx Virtex-6 FPGA (LX240T-2 FPGA standard build - up to LX550T-2 FPGA possible)
  - IDT CPS-1848 Serial RapidIO switch
    - SRIO V2.1 at up to 20Gbps per port with multiple x4 ports to backplane and front panel
  - Dual mini-SAS, to FPGA and SRIO - flexible, high-speed cabled connectivity
  - Three front panel SFP+ optical interfaces configurable as CPRI, OBSAI, GigE, S-RIO
  - Upcoming Expansion for DSP daughter card

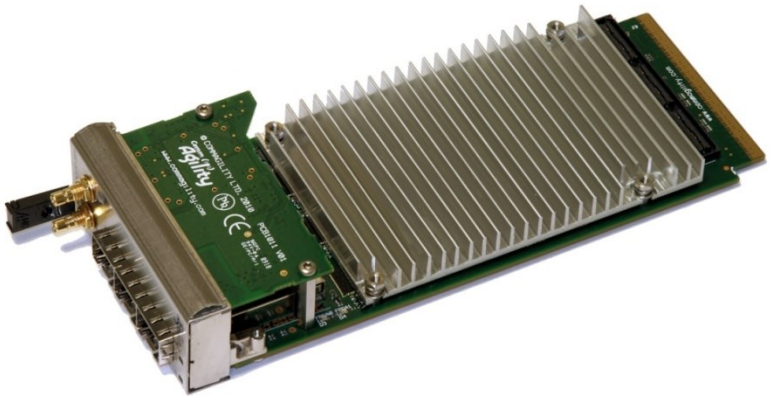

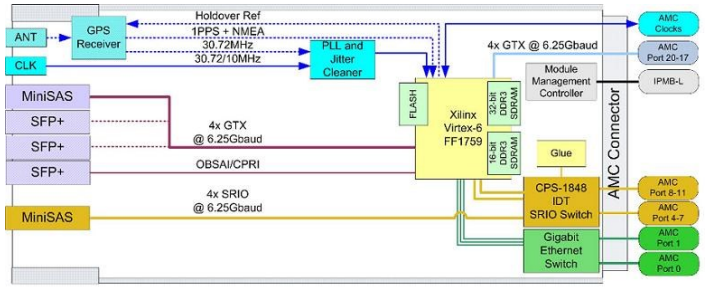

RapidIO Interconnect at CERN

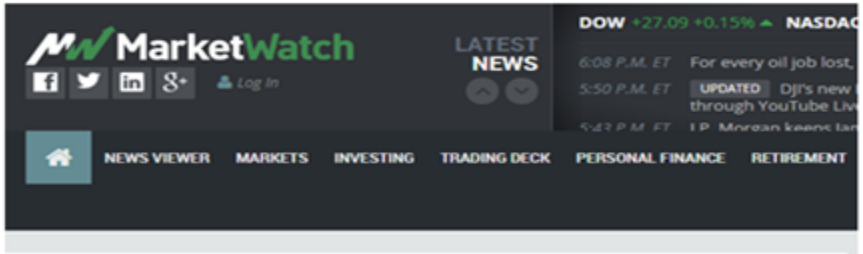

PRESS RELEASE

IDT Collaborates With CERN to Speed and Improve Data Analytics at Large Hadron Collider and Data Center

- Low latency interconnect fabric
- Heterogeneous computing
- Large scalable multi processor systems
- Desire to leverage multi core x86, ARM, GPU, FPGA, DSP with uniform fabric
- Desire programmable upgrades during operations before shut downs

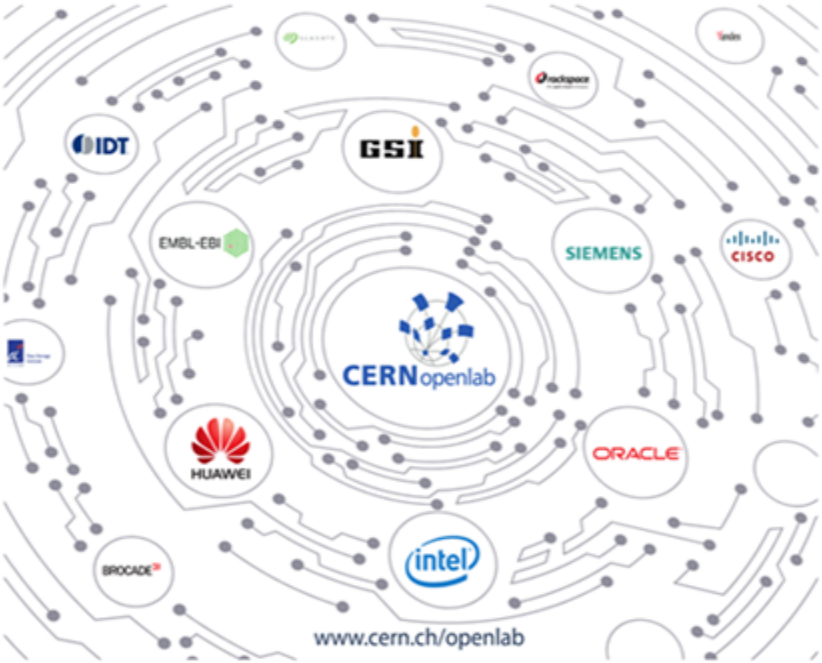

Edge Social Data Analytics

- Analyze User Impressions on World Cup Final 2014 (Germany/Argentina)
  - HPAC Lab project to analyze World Cup 2014 twitter data using Hadoop and visualize using Tableau public on HPAC Platform

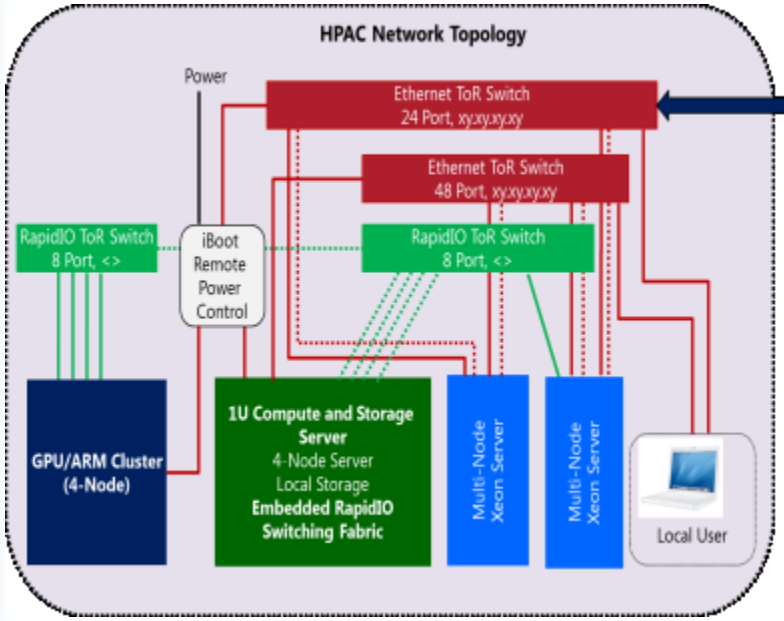

orange™ Silicon Valley

IDT RapidIO Product Portfolio and Roadmap

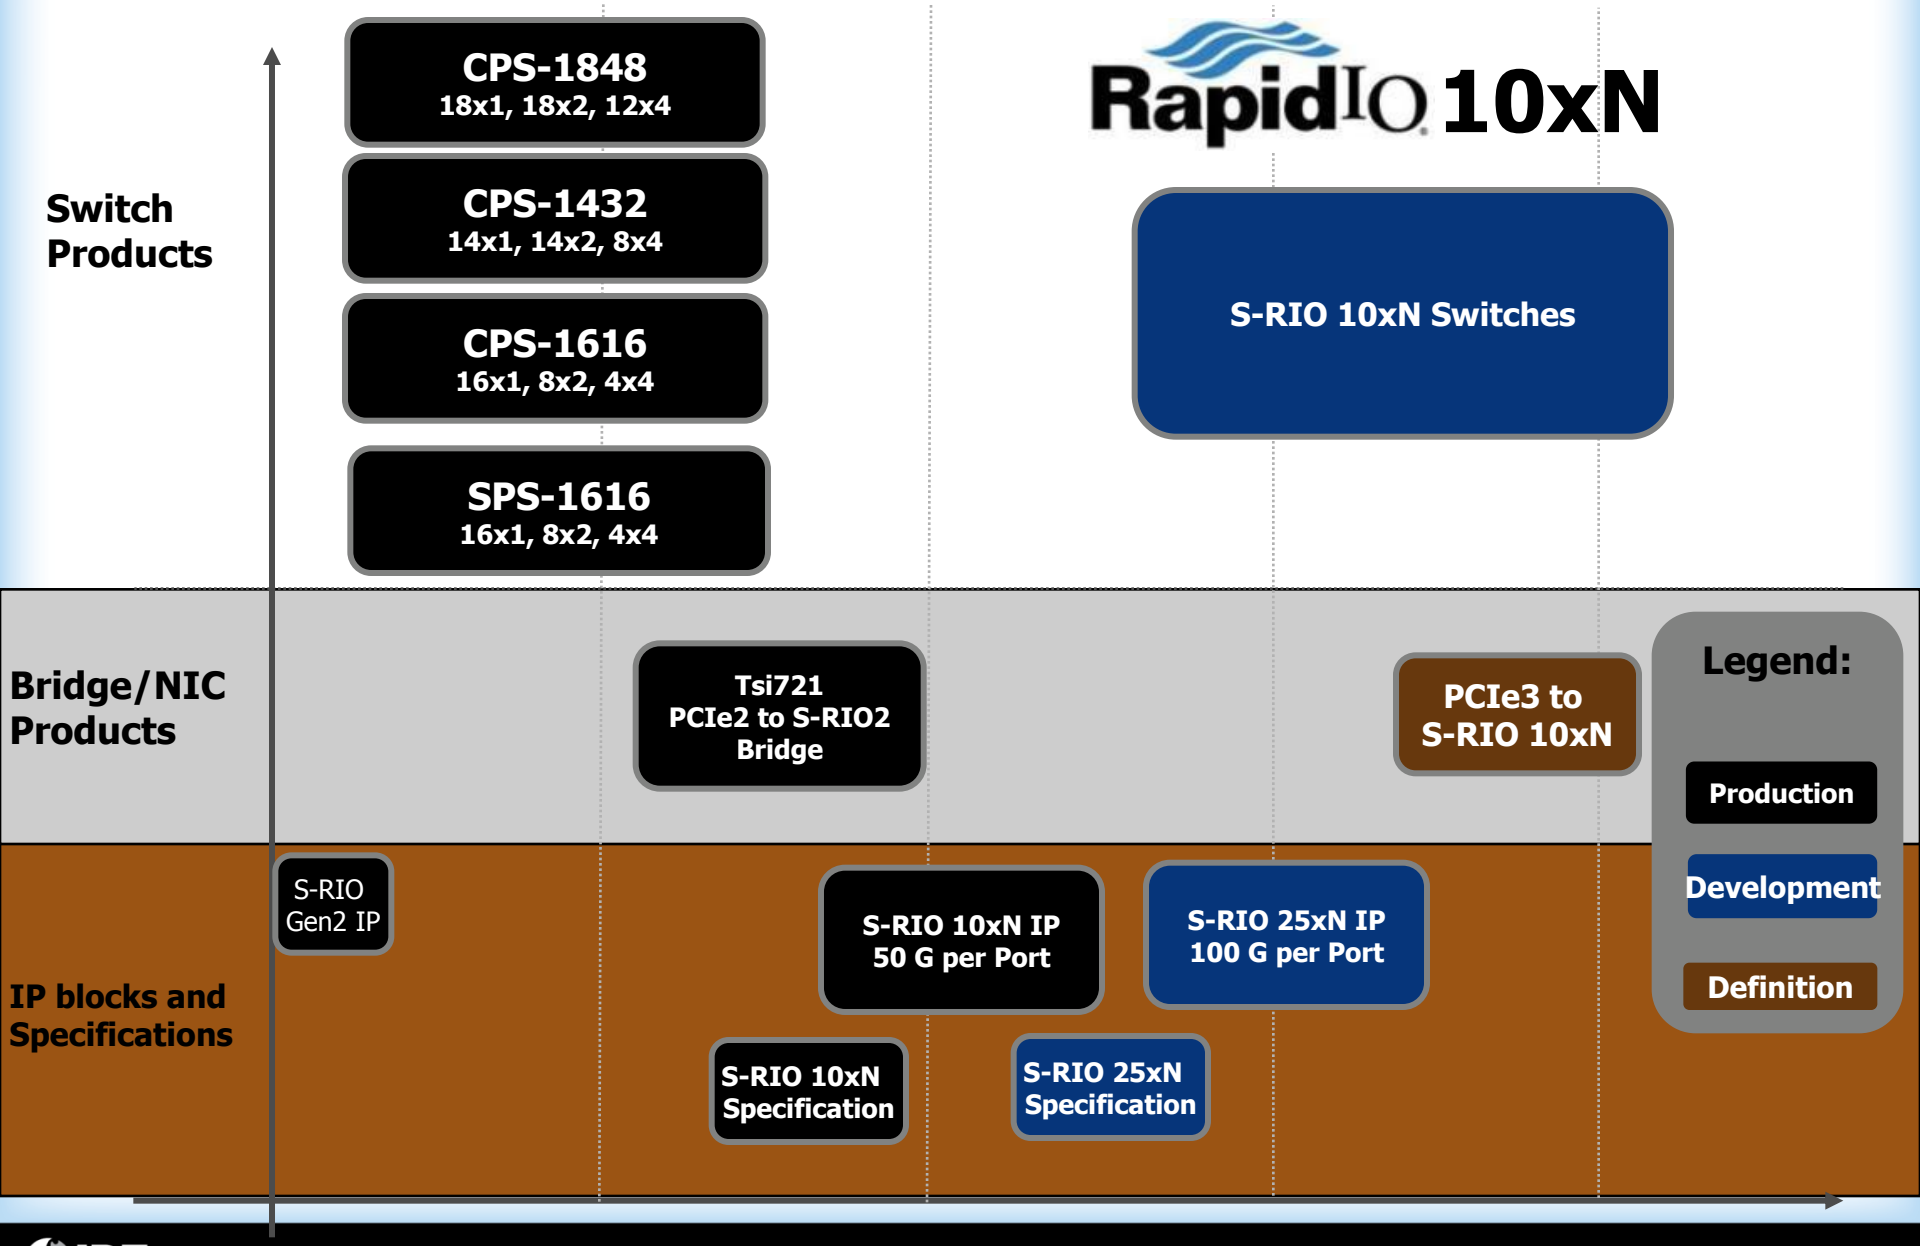

Summary of RapidIO 10xN in xTCA

Ideal for Embedded Mission Critical Systems

Low latency RapidIO provides a scalable low latency clustering fabric

x86, FPGA, DSP, PPC, ARM, GPU all supported in MicroTCA ecosystem

RapidIO 10xN MCH with 50 Gbps per port, 100 Gbps per AMC Interconnect in development for 2016

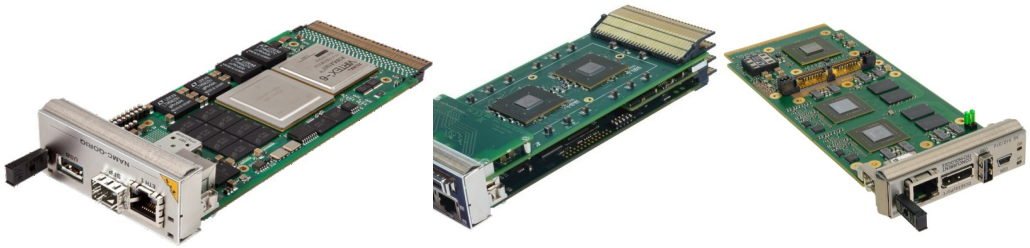

Development Tools

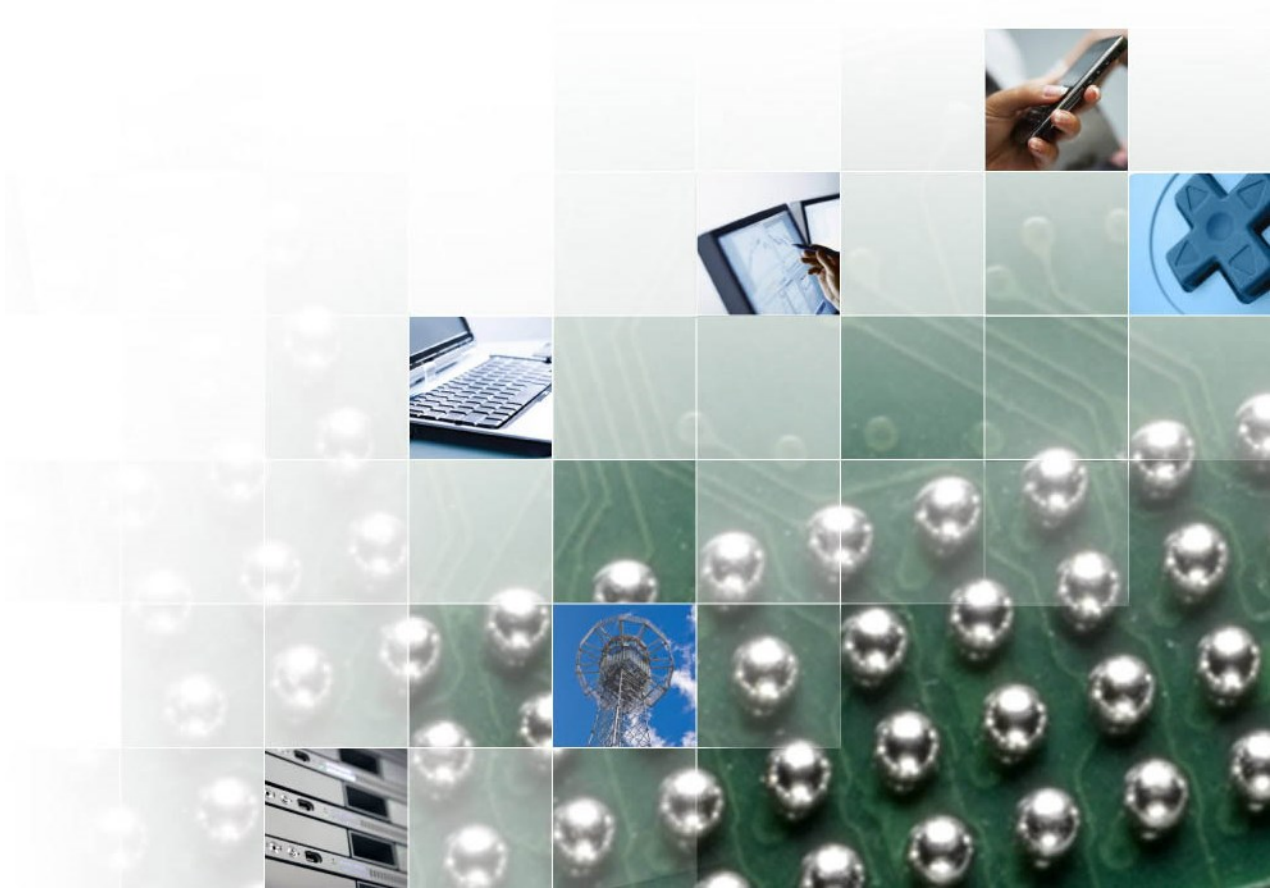

IDT S-RIO System Simulation Tool

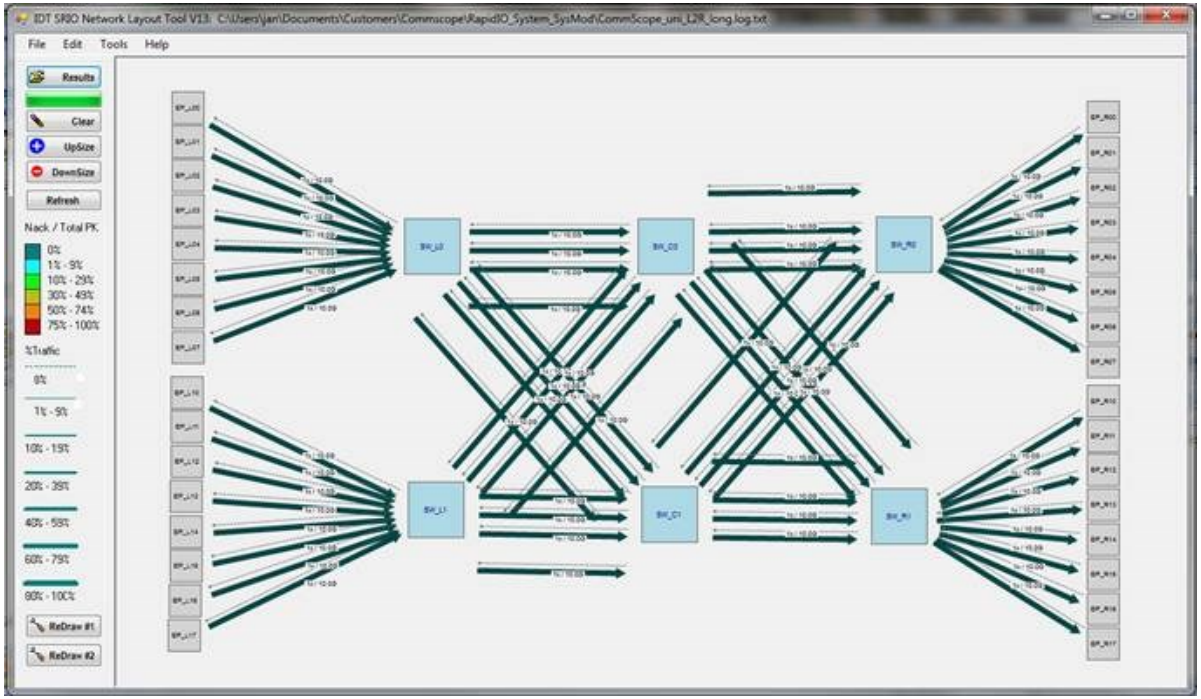

- Predict traffic throughput
- Useful for predicting end-to-end latencies for hi and lo priority packets
- Optimize link speeds and widths

Complex System Simulation

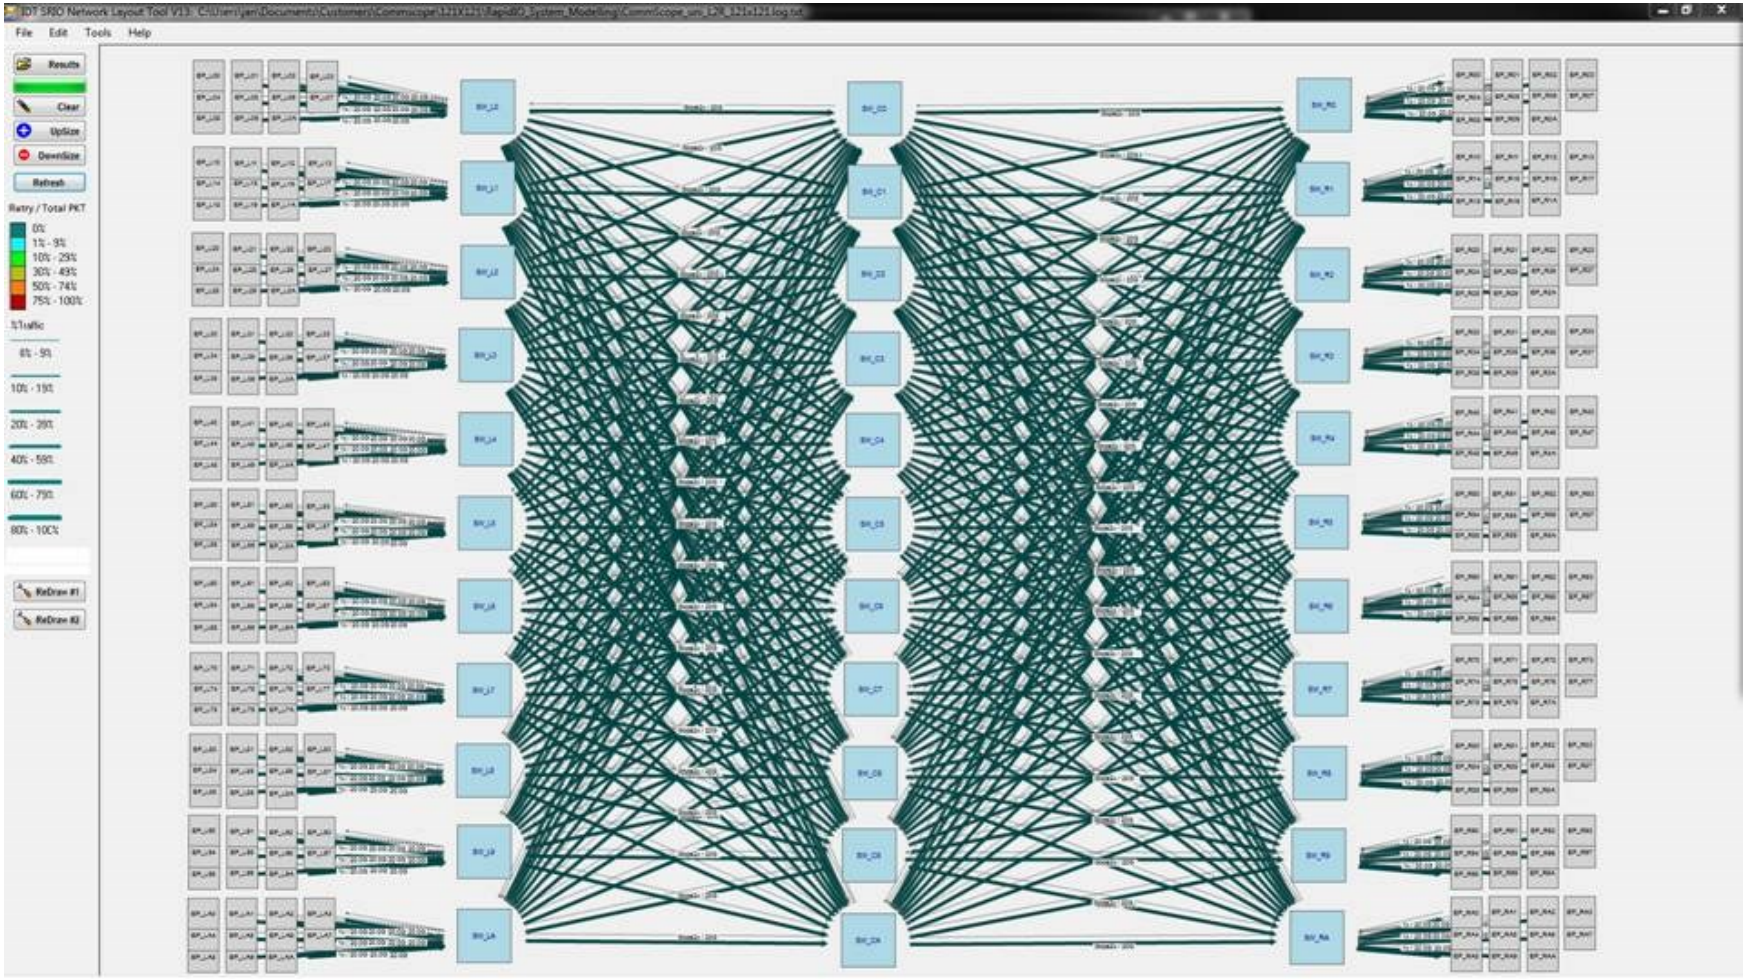

HW Debug Capabilities

- Event logs
- Packet capture
- BERT – On die scope function simplifies testing
- PRBS generation and checking
- Performance & diagnostic counters

S-RIO 10xN Development Platforms

Development Hardware

- Stand-alone platform that leverages widely recognized connectors throughout the industry

Development Software

- Register Access Tool
  - Accessible via I2C, JTAG and inband sRIO
  - Read and Write registers
  - Status commands.
  - Error log setup, and analysis
  - LUT programming and debug.
  - Packet generation and error injection control.
  - PRBS programming and usage
- On-die Scope Function
  - Integrated into Register Access Tool for easy register configuration and status reporting for easy design-in
  - Scope tool displays bit error curves

Used by OEM's and semiconductor suppliers for:

- Proof-of-concept system
- Interoperability testing
- Demonstration system
- Software Development
- ASIC/SOC/IP Development

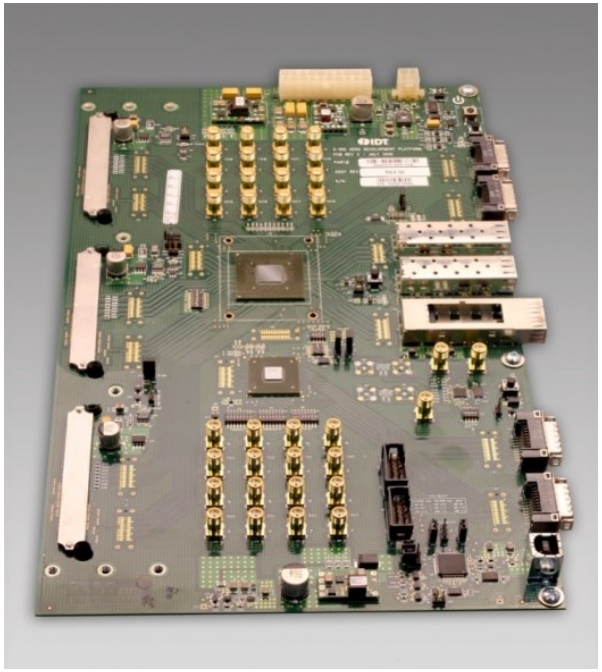

Open High Performance Analytics and Computing Lab

- Open Collaboration Lab around Heterogeneous Computing with low latency scalable RapidIO interconnect to accelerate end market usage
- Key focus areas:
  - Hyperscale Cloud Data Center-Based Analytics
  - High-Performance Computing
  - Wireless 4G advanced,
  - Wireless 5G and Mobile Edge Computing
  - Video Analytics
- Project by project contribution model
- Past, current and potential Projects
  - Twitter Analytics on FIFA World Cup Finals
  - Supercomputing at the Edge with GPU
  - CERN LHC target acquisition and data center analytics
  - Mobile Edge Computing for wireless networks
  - Edge Analytics for Financial Transactions and Fraud Detection
  - RapidIO ToR switching for Cloud RAN
  - Scalable low latency storage

Accelerate Adoption of OCP and RapidIO Solutions