CommAgility and N.A.T. CERN/HPC workshop

Flexible High Performance Architectures Based on MicroTCA.4 and RapidIO

Edward Young / Vollrath Dirksen

June 2015

edward.young@commagility.com

vollrath@nateurope.com

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

What do we need for HPC?

- Heterogenous Computing
  - ARM, DSP, FPGA, GPU etc
- Suitable Interconnect
  - low latency, efficient, endpoints, speed
  - Plus Ethernet (of course)
- Performance Density and reasonable cost
- System Management
- Suitable software architectures
- What about compatibility with other systems?

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

Our Proposal

- MicroTCA.4
  - Used for front end data capture in Physics experiments
  - This already contains the elements to build HPC
- Add Serial RapidIO
  - Currently Gen2 with all hardware available now
  - Standard option for MCH, dual star network to cards
  - Standardized software protocols for communication
- A range of processing options, e.g.
  - ARM+DSP with FPGA
  - High end FPGA
- Summary

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

MTCA.4 µRTM Extension

µRTM fully managed by front AMC

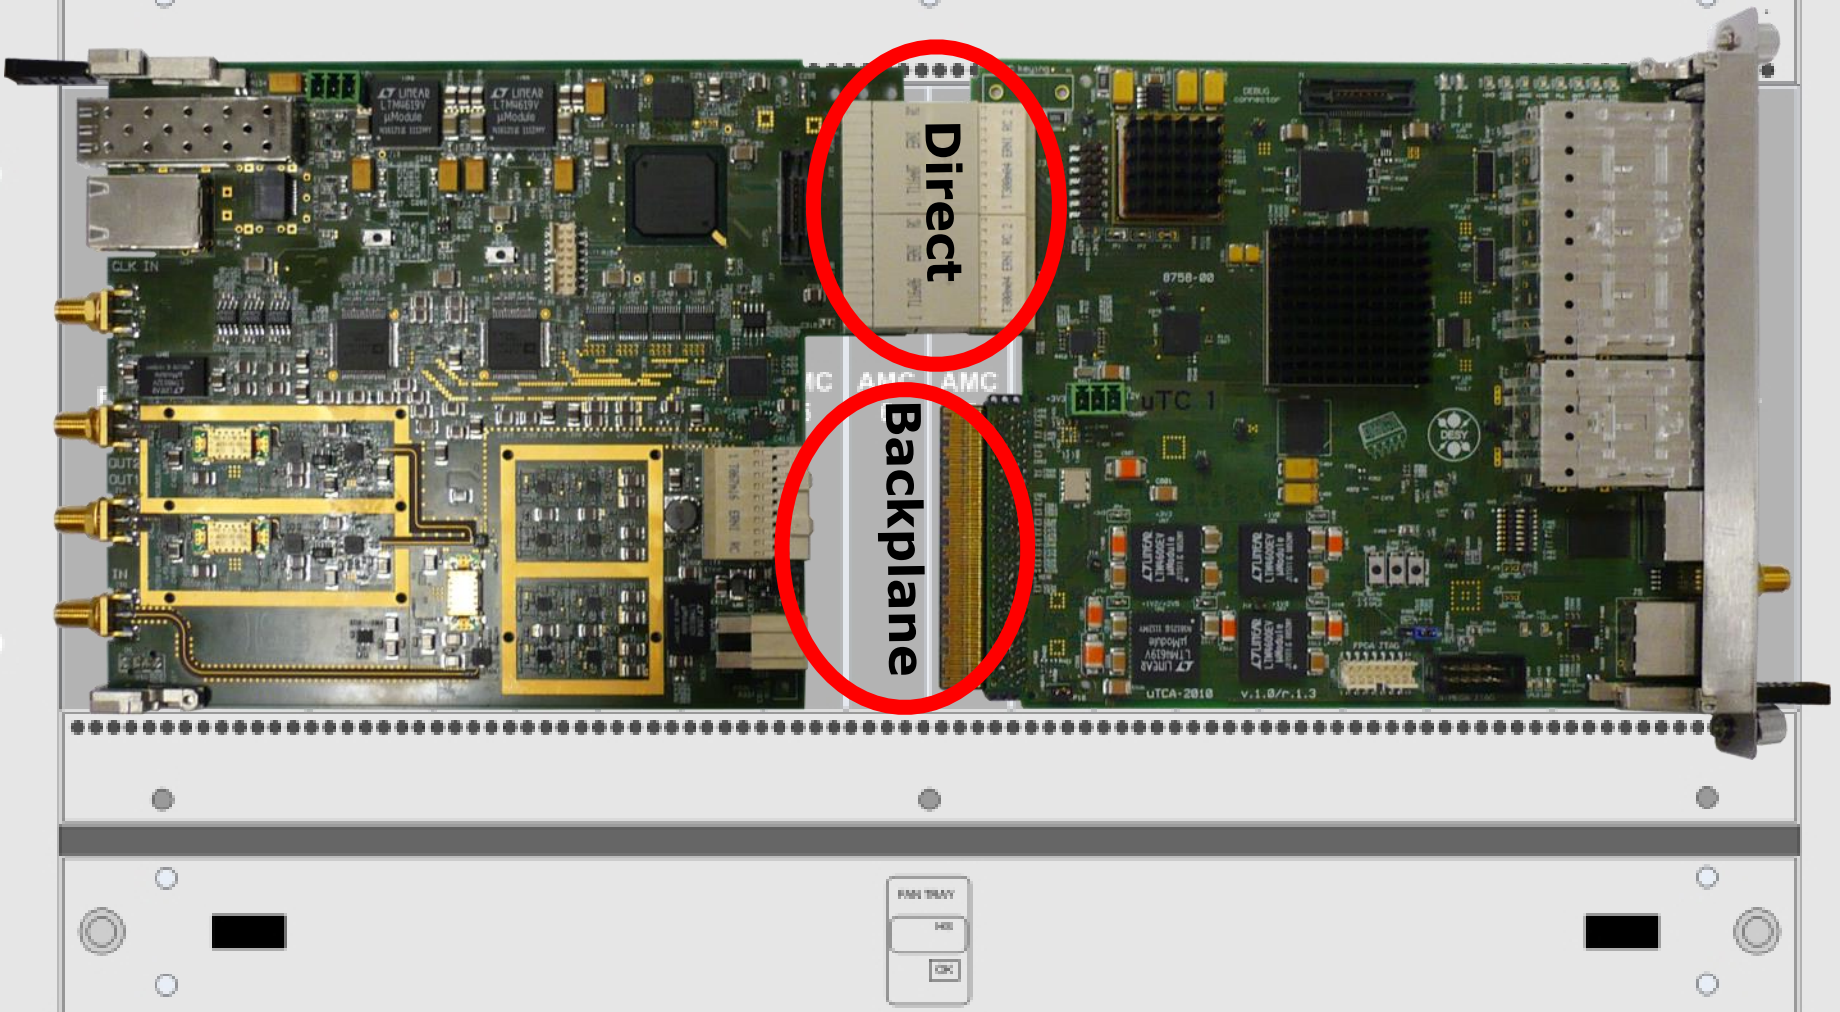

MTCA.4 Chassis

19'', 7U, 12-Slot

12913800

or below

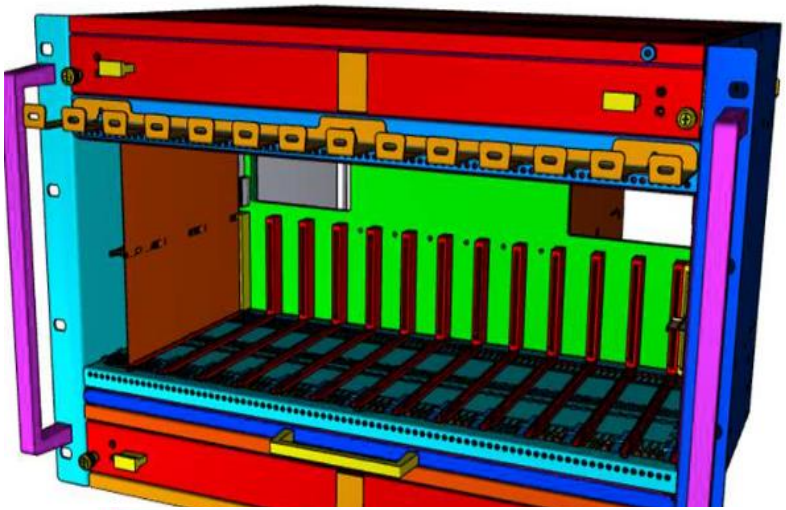

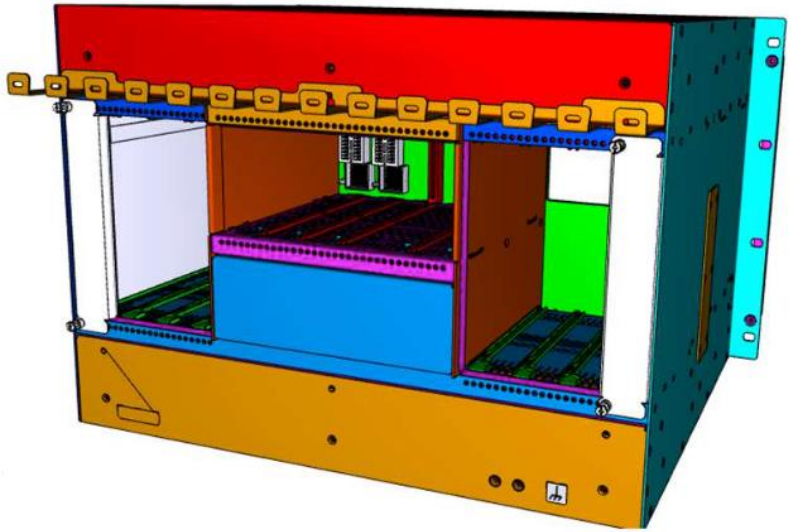

7	Rear card cage	9	Ground Terminal
8	Cable Tray (Can be mounted above or below the card cage)	5	Card Cage with Guide Rails

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

MicroTCA Architecture
MTCA backplane – multiple transfers

- GbE
- SATA
- fat pipe
  - PCIEx4
  - SRIOx4
  - XAUI
- Clock /USB
- JTAG
- IPMI

MCH1 AMC 1-12 MCH2

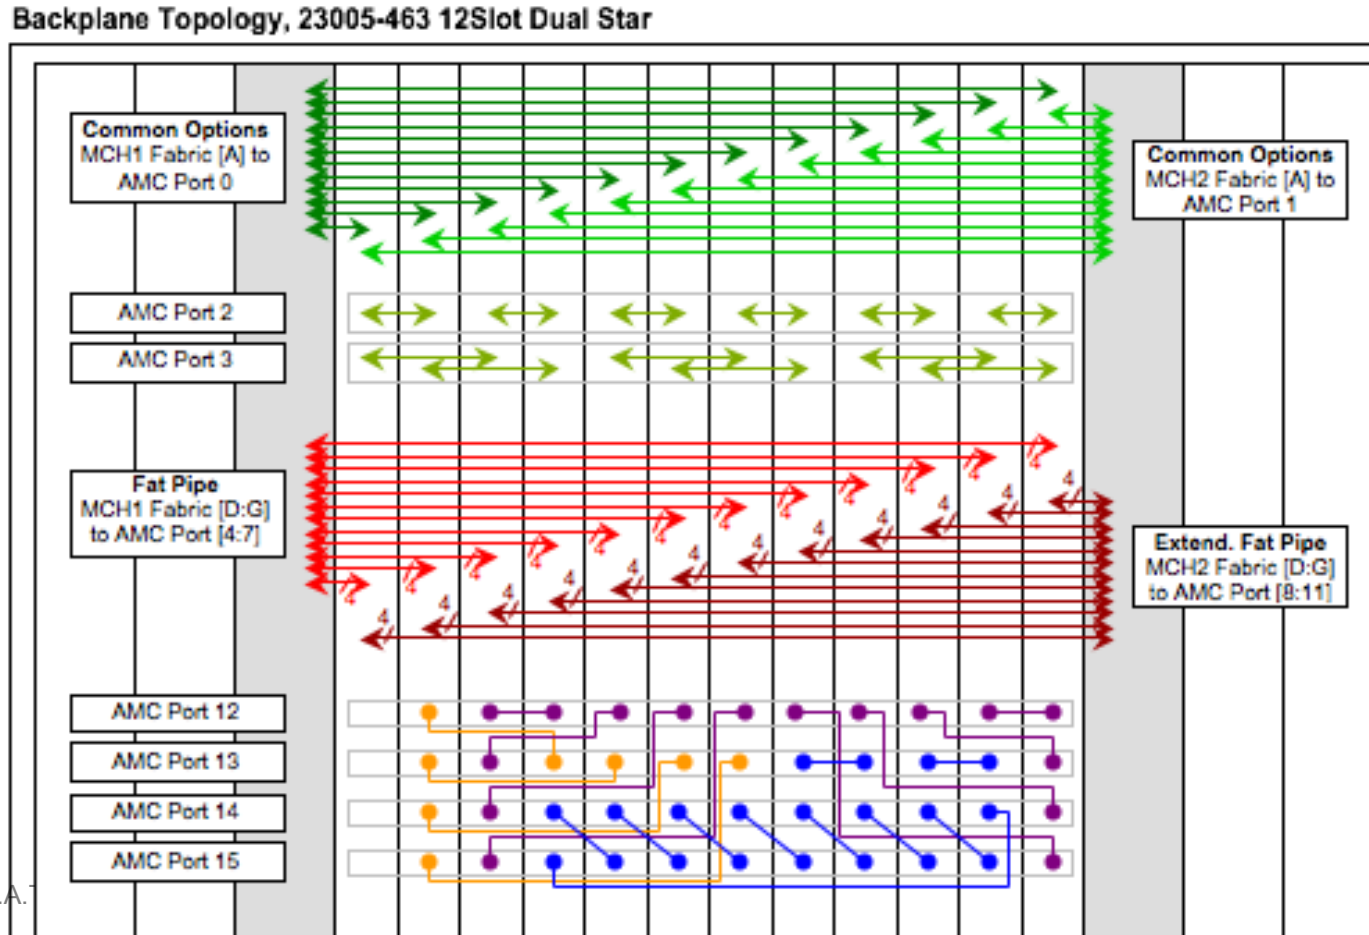

MicroTCA
Architectural features

- simple backplane architecture
  - reduces costs and risks, is re-useable in future
- all signals at same signal level (MLVDS)
  - no electrical clash
- switched connections
  - no blocking transfer
  - type of backplane connection depends on kind of switch
- all slots managed and controlled
  - detection of incompatibilities and faults
  - health management and fault isolation
  - hot-swap and hot-plug

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

MicroTCA Architecture
Infrastructure component: power modules

- NAT-PM-DC420 Input DC -48V Payload: 420W
- NAT-PM-DC840 Input DC -48V Payload: 840W
- NAT-PM-AC600 Input AC 110-265 Payload: 600W
- NAT-PM-AC600D Input AC 110-265V Payload: 600W (double width)
- NAT-PM-AC1000 Input AC 110-265V Payload: 1000W (double width)
- NAT-RPM-PSC Input AC 110-265V Payload: 600W (double width)

- Features:
  - monitoring of all 16 power channels
  - load sharing
  - n+1 redundancy
  - load bar

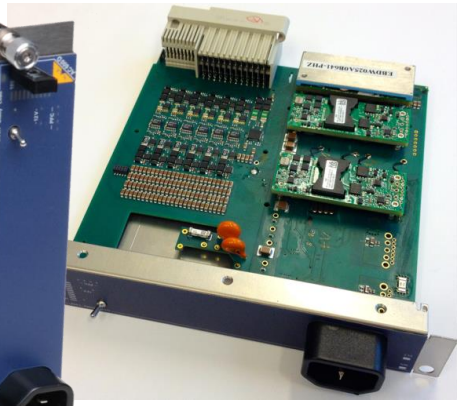

DC420 DC840 AC600 AC600D AC1000 (2015) RearPM for LLRF

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

JTAG-Switch-Module
NAT-JSM

- MTCA.x compliant JTAG Switch Module
- Any AMC, CU, PM

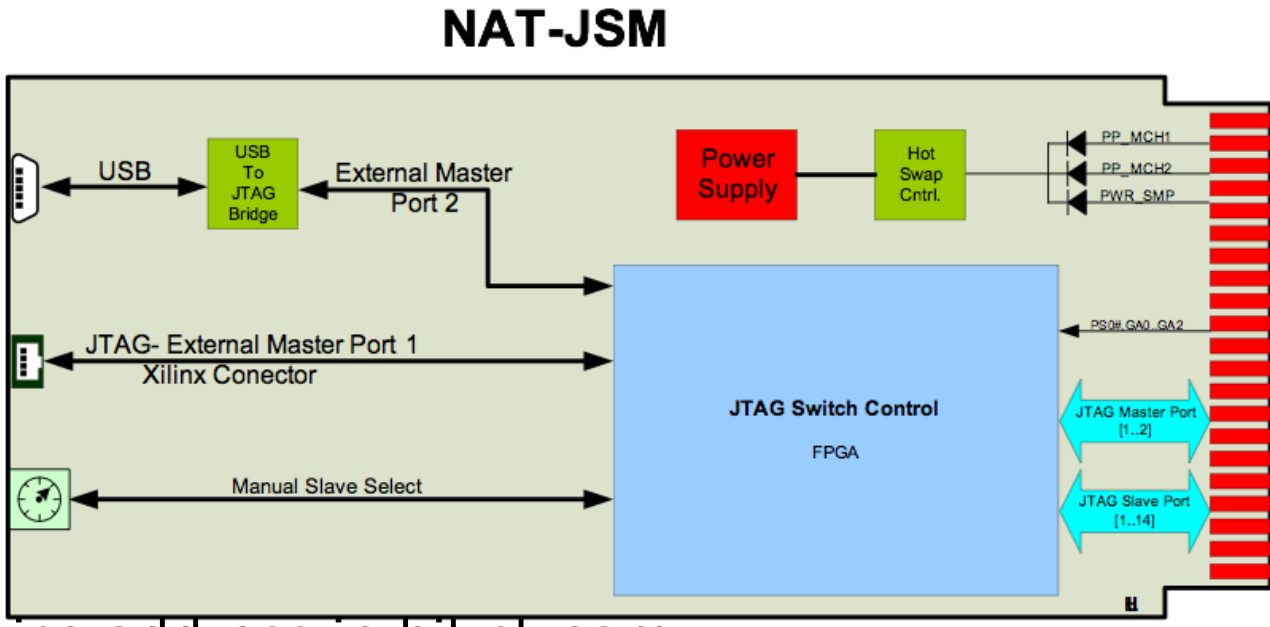

- Programming:
  - Xilinx Connector
  - USB
  - **MCH**
- Target selection by rotary switch or device address in bit stream
- Master auto detection

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

Our Proposal

- MicroTCA.4
  - Used for front end data capture in Physics experiments
  - This already contains the elements to build HPC
- Add Serial RapidIO
  - Currently Gen2 with all hardware available now
  - Standard option for MCH, dual star network to cards
  - Standardized software protocols for communication
- A range of processing options, e.g.
  - ARM+DSP with FPGA
  - High end FPGA
- Summary

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

NAT-MCH: Single & Double

Base, CLK, Fatpipe (PCIe, XAUI, SRIO), Custom

NAT-MCH

SRIO-Submodule: Block-Diagram

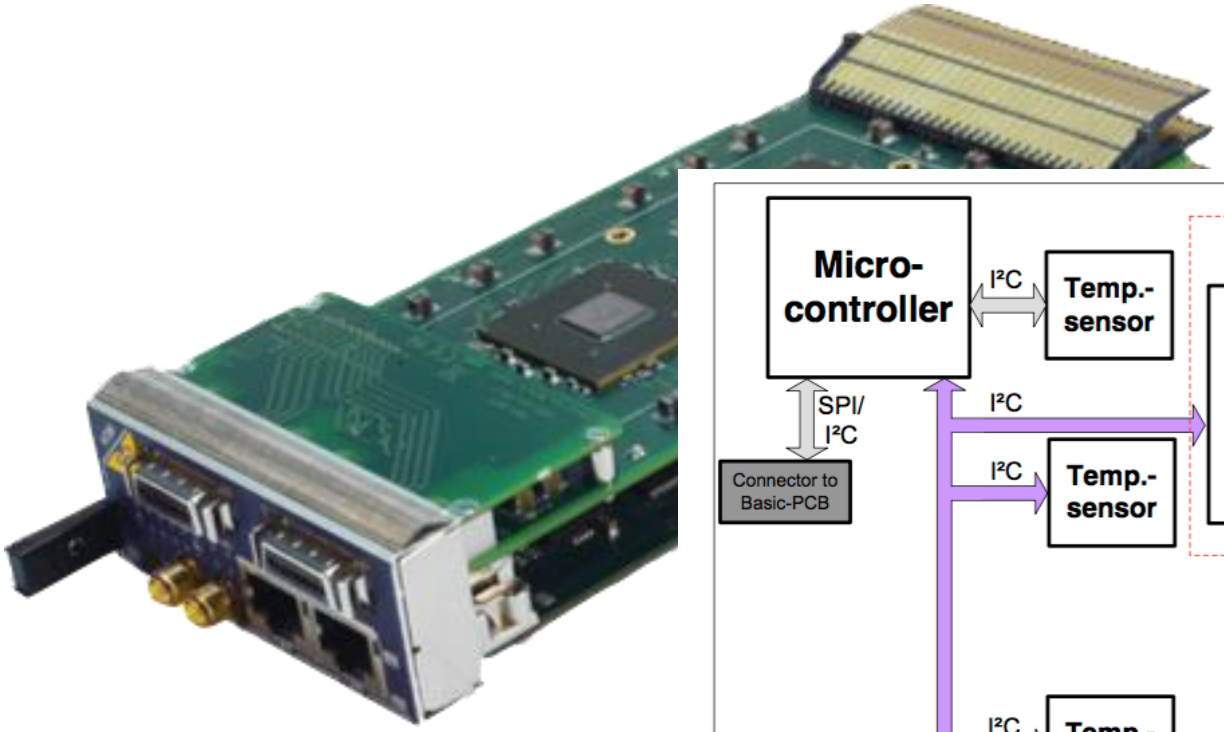

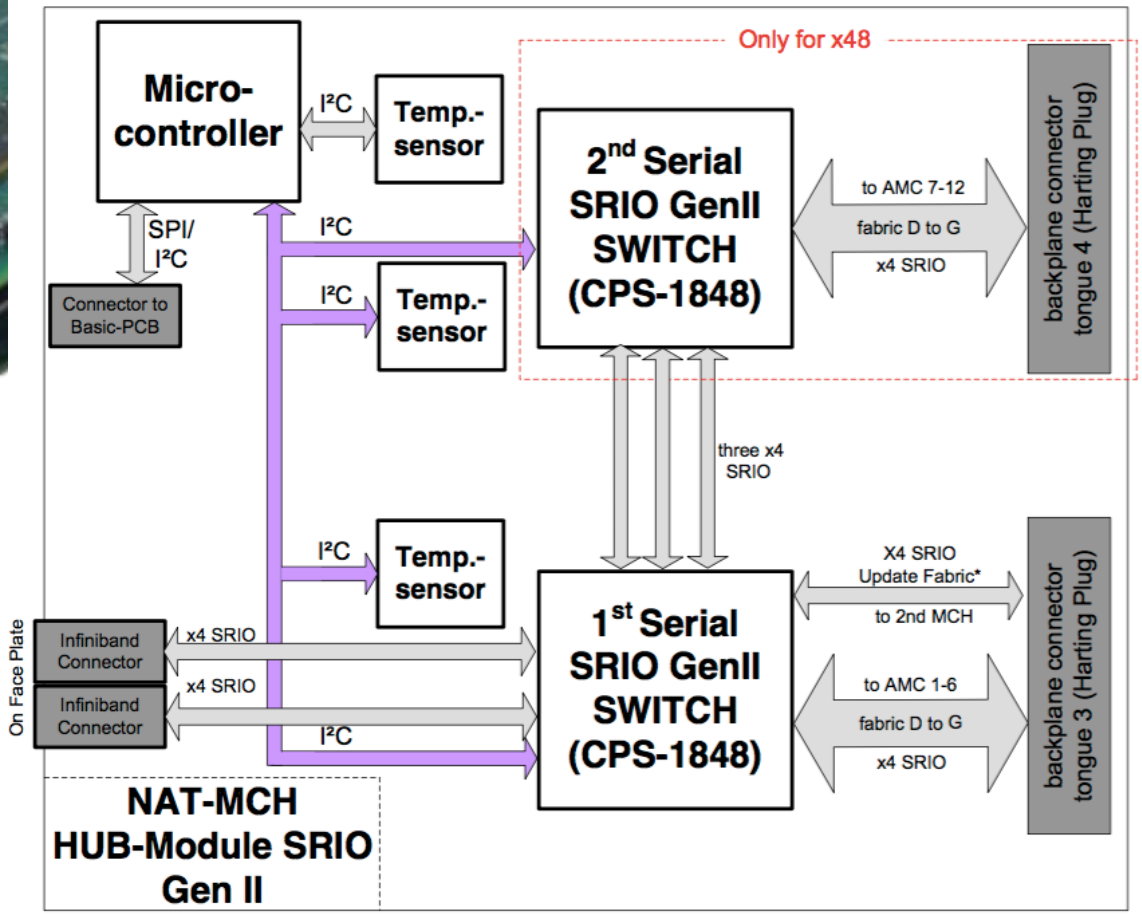

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

MTCA.4 with 12 AMCs and 2 MCHs

Redundant SRIO connections

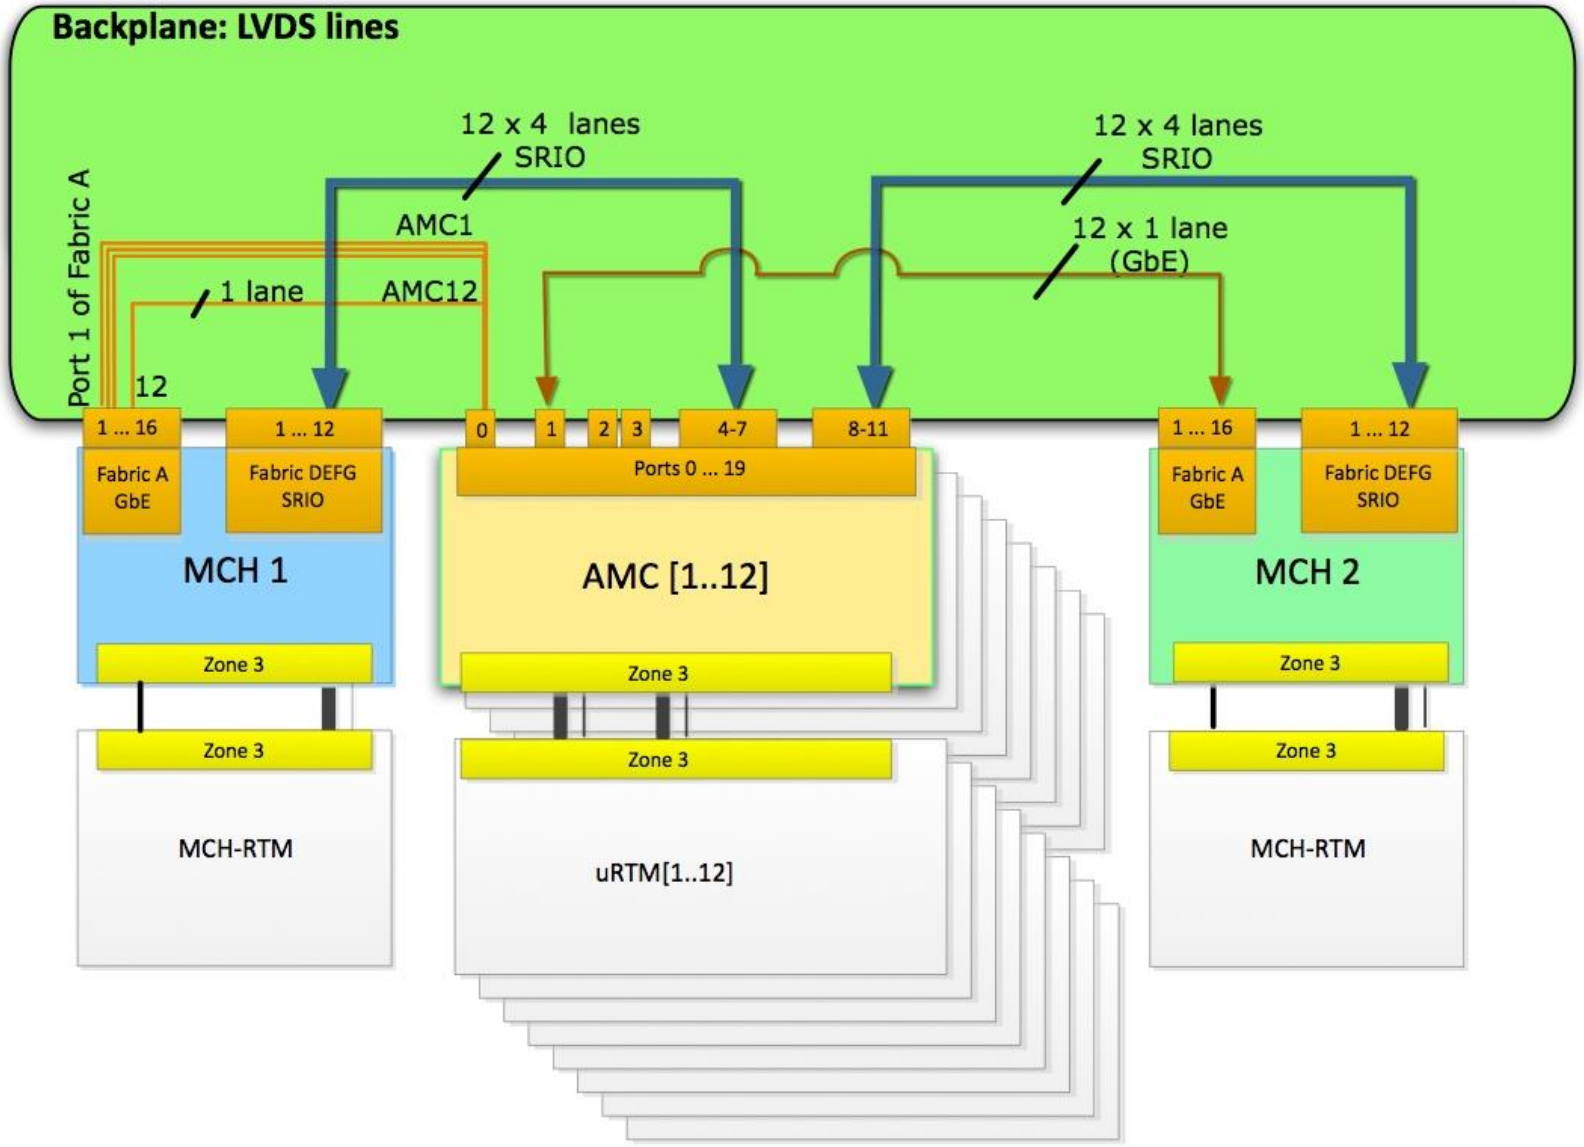

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

Our Proposal

- MicroTCA.4
  - Used for front end data capture in Physics experiments
  - This already contains the elements to build HPC
- Add Serial RapidIO
  - Currently Gen2 with all hardware available now
  - Standard option for MCH, dual star network to cards
  - Standardized software protocols for communication
- A range of processing options, e.g.
  - ARM+DSP with FPGA
  - High end FPGA
- Summary

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

ARM+DSP+FPGA: AMC-D24A4

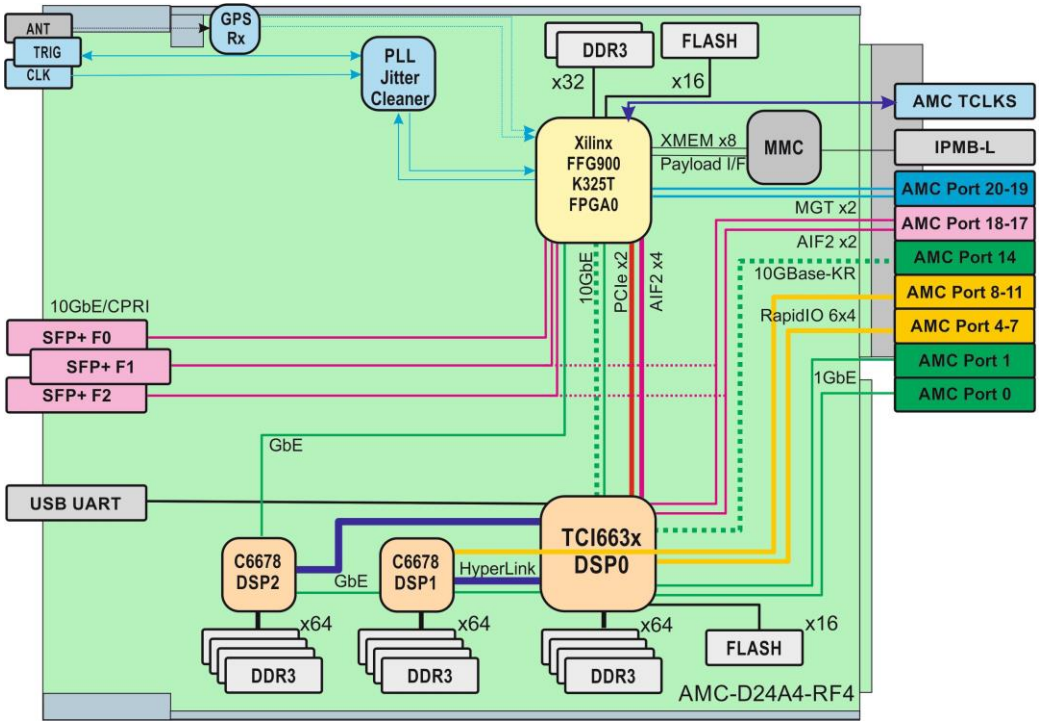

- ARM: 4 x A15 cores @ 1.4GHz
- DSP: 24 x C66x cores @ 1.2/1.25 GHz, plus built in accelerators (e.g. FFT)
- FPGA: Kintex-7 K325T with local PCIe to main SoC for acceleration
- Serial RapidIO: 2 off 6x4 20Gbps to backplane
- Other I/O: optical, GbE, timing, and 10GbE possible
- Mid-size AMC; Mezzanines possible if full-size
- Available now

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders Confidential

DSP Platform Support

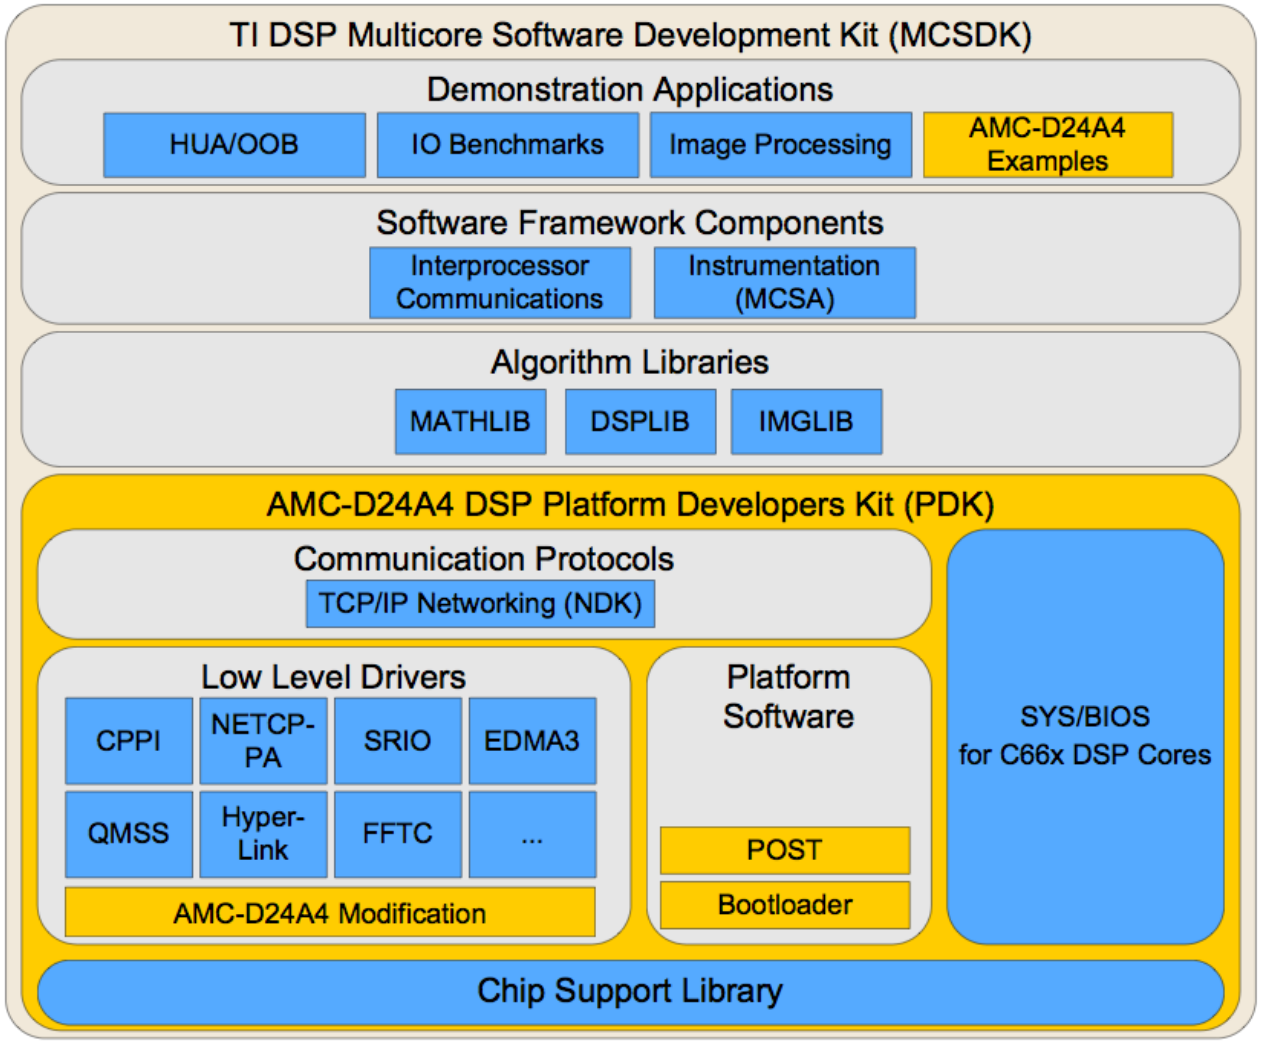

- OpenCL and OpenMP support available
- Open SRIO drivers

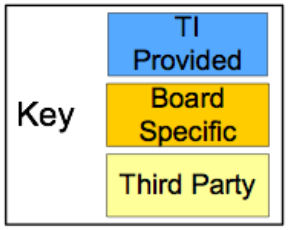

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders Confidential

Linux Platform Support

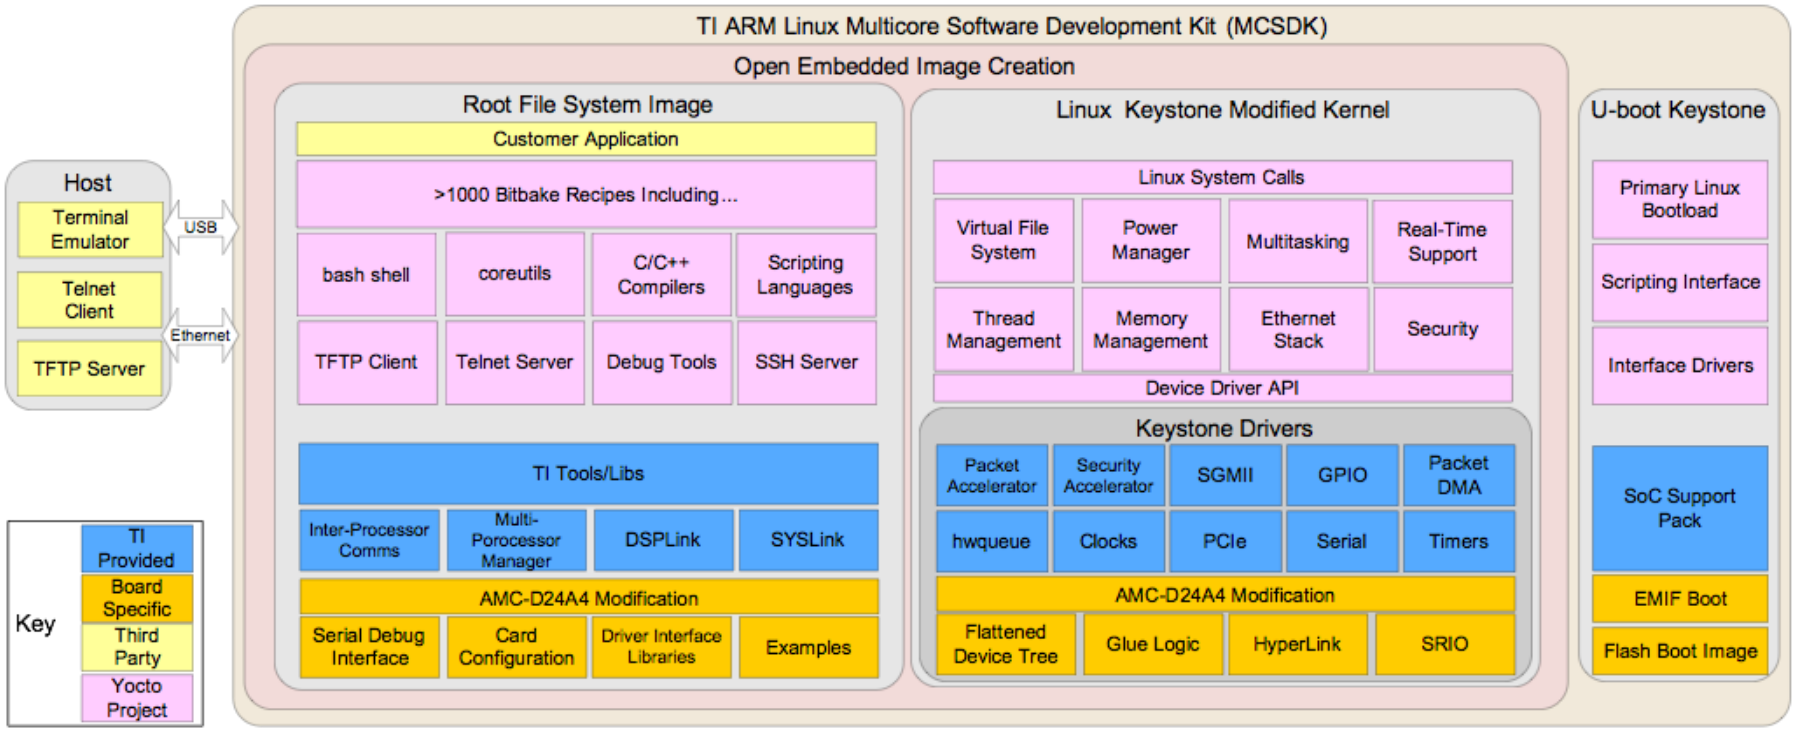

FPGA Support

- Provides tools to diagnose and reprogram FPGA
- Examples provided to exercise card interfaces

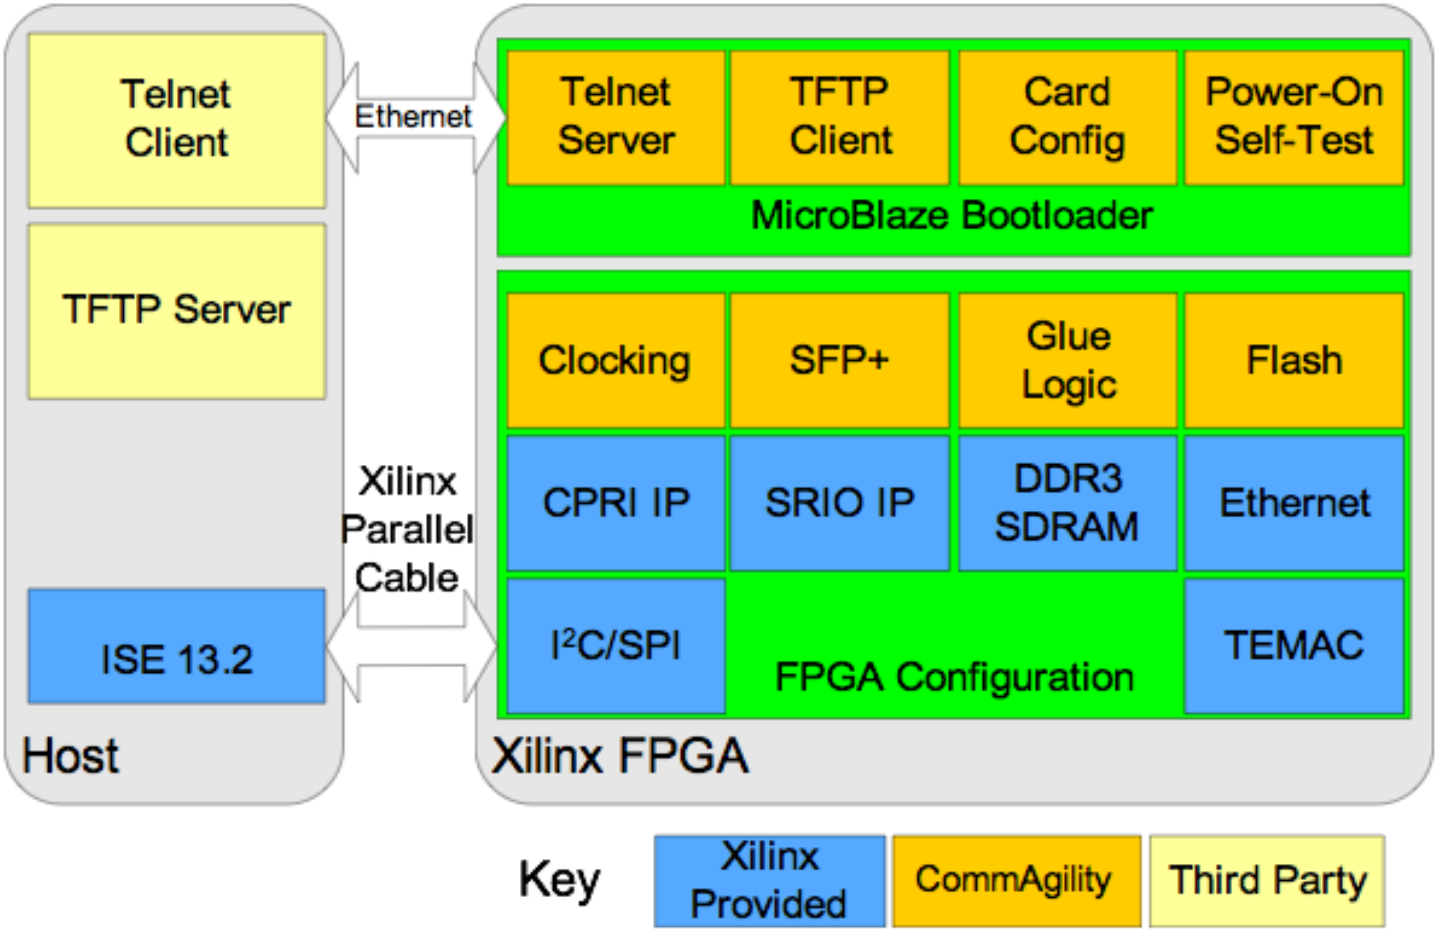

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders Confidential

Our Proposal

- MicroTCA.4
  - Used for front end data capture in Physics experiments
  - This already contains the elements to build HPC
- Add Serial RapidIO
  - Currently Gen2 with all hardware available now
  - Standard option for MCH, dual star network to cards
  - Standardized software protocols for communication
- A range of processing options, e.g.
  - ARM+DSP with FPGA
  - High end FPGA
- Summary

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

FPGA Processing Card

- Roadmap product for CommAgility, currently in discussion with a lead customer for 2016
- An extremely high performance and flexible card for HPC
- FPGA into HPC and data centres is going mainstream, but this system can offer much better interconnect between FPGAs
- FPGA tools and C based synthesis are improving all the time

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

FPGA Card Block Diagram

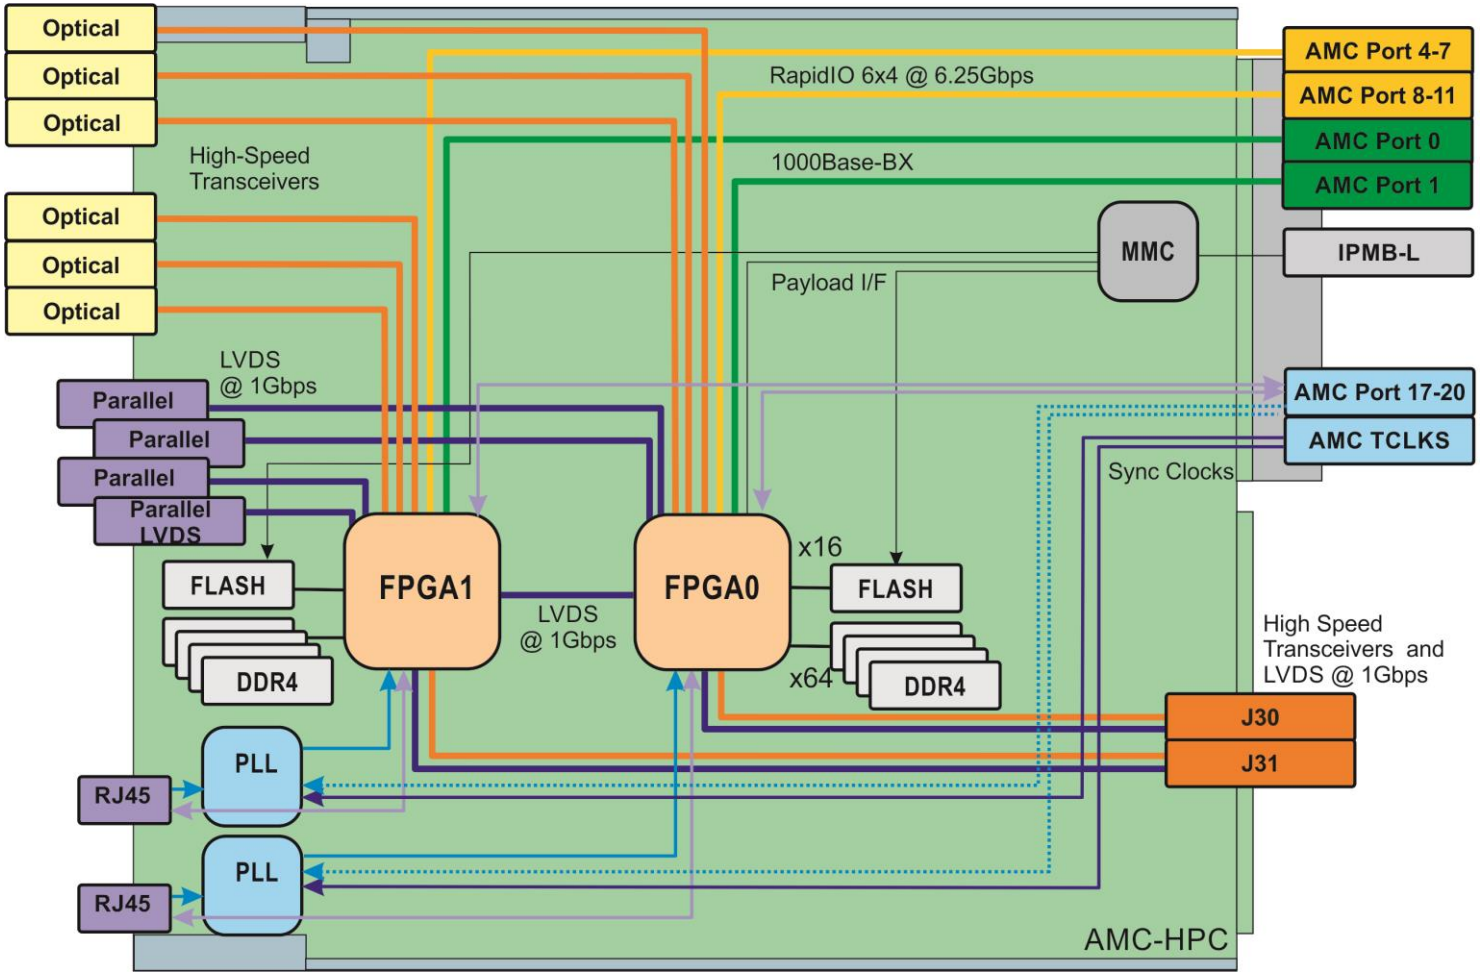

- Two high end FPGAs (Ultrascale / Stratix 10): Approx 10M Logic cells per board, plus DSP, logic, memory etc

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

Our Proposal

- MicroTCA.4
  - Used for front end data capture in Physics experiments
  - This already contains the elements to build HPC
- Add Serial RapidIO
  - Currently Gen2 with all hardware available now
  - Standard option for MCH, dual star network to cards
  - Standardized software protocols for communication
- A range of processing options, e.g.
  - ARM+DSP with FPGA
  - High end FPGA
- Summary

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders Confidential

Performance Estimates

- Let's consider a full MTCA.4 chassis
  - 19" rack mount, 7U high, 12 processing cards
  - Fully managed, reliable, front to rear cooling
- Filled with DSP/ARM cards
  - ~5.6 TFLOPS and 11.2 TMACS from DSP cores
  - ~210 Dhrystone MIPS from ARM cores
  - ~4M FPGA Logic cells plus 10K DSP slices
- Filled with FPGA cards
  - ~120M FPGA Logic cells plus 50K DSP slices
- All with 40Gbps of efficient, low latency Serial RapidIO connection to each card
- Any combination of the above, plus 3rd party cards
- Local HPC could be in same chassis as data acquisition

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

Summary: What do we need for HPC?

- Heterogenous Computing
  - We've shown ARM, DSP, FPGA. Standard based allows others
- Suitable Interconnect
  - Serial RapidIO
- Performance Density and reasonable cost
  - An efficiently packaged smaller system
- System Management
  - Good management and reliability is inherent in MTCA
- Suitable software architectures
  - Standardised interconnect protocols
  - Supporting OpenCL etc
- What about compatibility with other systems?
  - Links in with MTCA.4 data acquisition systems

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders Confidential

ONE TECHNOLOGY
MULTIPLE SOLUTIONS

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders

Thank you very much!

Questions?

Edward Young
edward.young@commagility.com

Commagility
Charnwood Building Holywell
Park Ashby Road
Leicestershire Loughborough
www.commagility.com

Vollrath Dirksen
vollrath@nateurope.com

N.A.T. GmbH
Konrad-Zuse-Platz 9
53227 Bonn, Germany
www.nateurope.com

N.A.T. Training:

MTCA.4 Basic
MTCA.4 Advanced

Coming soon: Webinars

 © 2015 CommAgility & N.A.T. GmbH I All trademarks and logos are property of their respective holders