Introduction to MicroTCA.4

ARD ST3 Workshop 2021

Cagil Gumus
Hamburg, 29.09.21

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

DESY.

Motivation

You

This tutorial

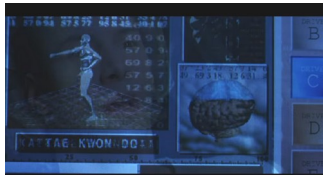

Your supervisor

The Matrix 1999

Things that will be covered..

What is MicroTCA?

Brief History/Evolution of MicroTCA

Building Blocks of MicroTCA.4 Systems:

- The crate
- AMC + RTM
- MCH
- Backplane
- RF Backplane

Management

Protocols on the backplane:

- PCIe
- Point to Point Links
- Timing Clocks and Interlocks (MVLDS)

What is MicroTCA

What is µTCA®

- It is an **open standard, embedded computing specification**.
- Specification is written by PICMG (PCI Industrial Computer Manufacturers) Group:
  - nonprofit consortium of companies and organizations
- Defines **fully** and **non-redudant** system configurations
  - Serviceability, Availability → **up to 99.9999% uptime (six-nines)** → **31.56 seconds downtime in a year**
- **Power budgetting** and **Hot Swap** Capabilities
- Comes with complete **component and system management** that **allow failure detection** and **isolation**

Family Tree of PICMG Specifications

- MicroTCA has evolved from ATCA
- Main motivation of MicroTCA:
  - ‘Lightweight’ version of ATCA
  - Plug the AMC cards directly on a backplane
- MicroTCA.0 is the base specification
- MicroTCA.1,2,3,4 are just **subsidiary specifications** on top of the base specification

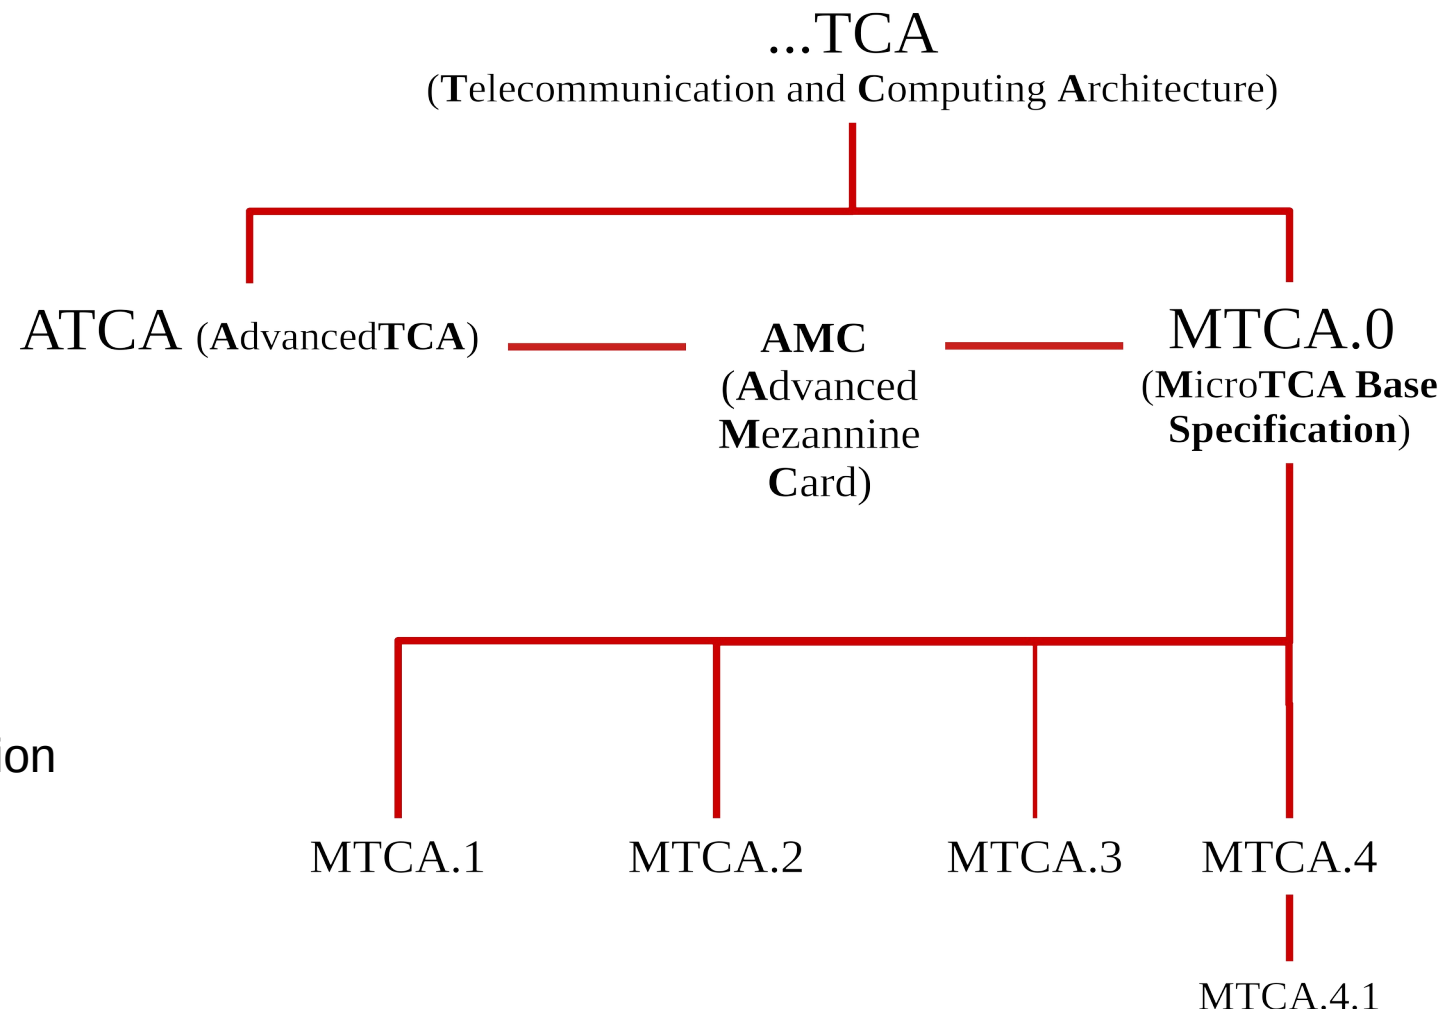

Family Tree of PICMG Specifications

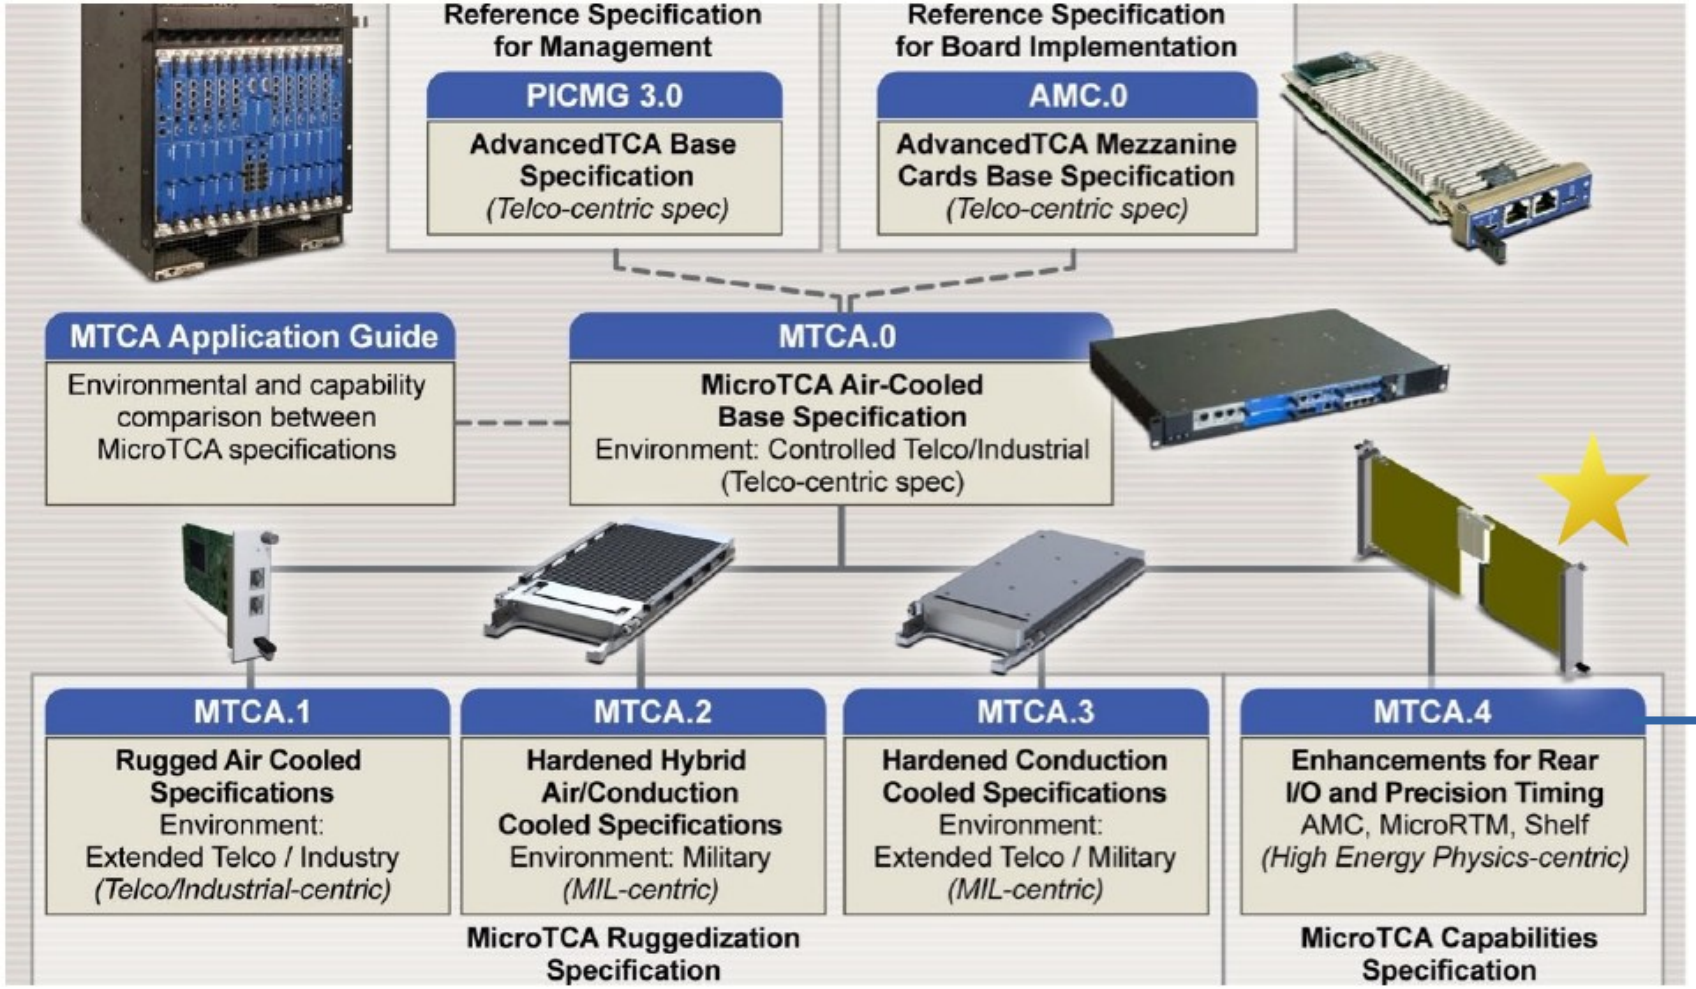

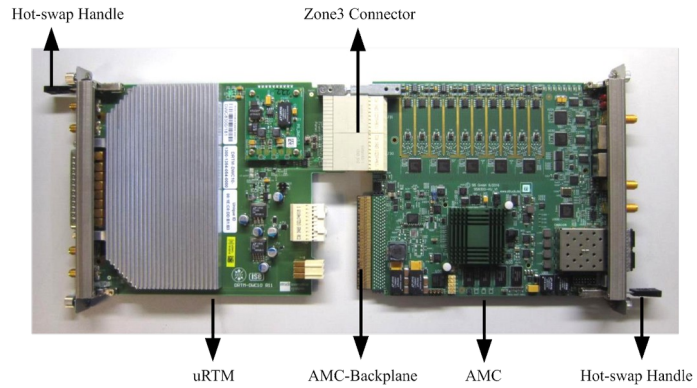

- **R**ear **T**ransition **M**odule (**RTM**) gets introduced
- RF Applications (eg. Experimental Physics)

Quick look at ATCA Systems

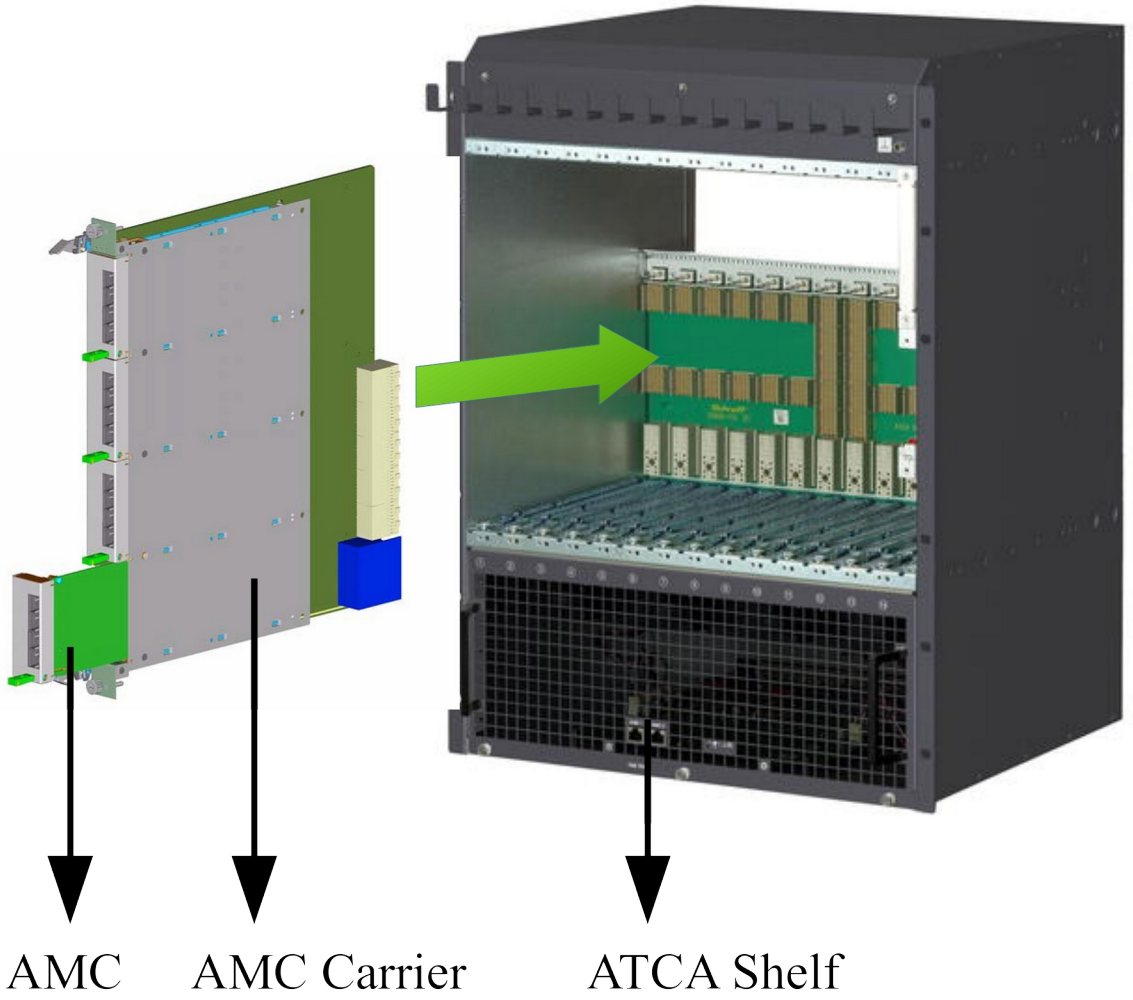

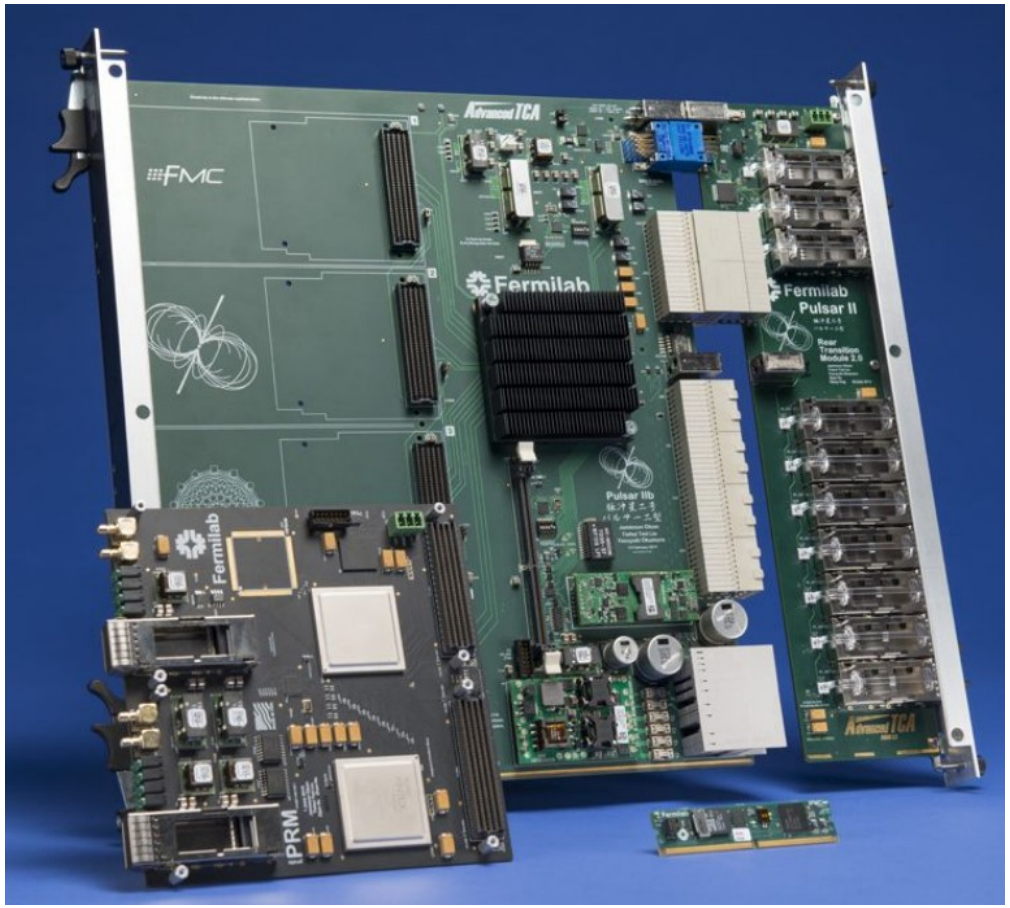

The Crate

MicroTCA Crates

The trade off between functionality – redundancy – reliability

1U MicroTCA.4

- Integrated eMCH
- 2x Double Mid-Size AMC slots w/ RTM
- 2x Single Mid-Size AMC slots
- Integrated 400W PSU
- No redudant part

3U MicroTCA.4

- 4x Double Mid-Size AMC slots w/ RTM
- 1x Double Mid-size AMC slot
- 1x Double Full-size AMC Slot
- Discreete PSU
- No redudant part

5U MicroTCA.4 (Cube)

- 6x Double Mid-Size AMC slots w/ RTM
- 1x Double Full-size MCH slot
- 1x Double Full-size AMC Slot
- Discreete PSU
- No redudant part

9U MicroTCA.4

- 6x Double Mid-Size AMC slots w/ RTM
- 1x Double Full-size MCH slot
- 1x Double Full-size AMC Slot
- Discreete PSU
- Redudant MCH + PM
- RF Backplane capable
- JTAG Switch Module available

AMC + RTM

AMC + RTM

- **AMC**:
  - Carries mostly the ‘digital logic’.
  - Many COTS components available
  - Can be single or double width
- **RTM**
  - Extention of the connected AMC
  - Big dependency on the application (in-house development)
  - RF Front ends, CPUs, GPIO …
- RTM is the subordinate of the AMC in terms of management
- AMC card gets connected to the crate via AMC-Backplane connector
- RTM and AMC connects through **Zone 3** connector

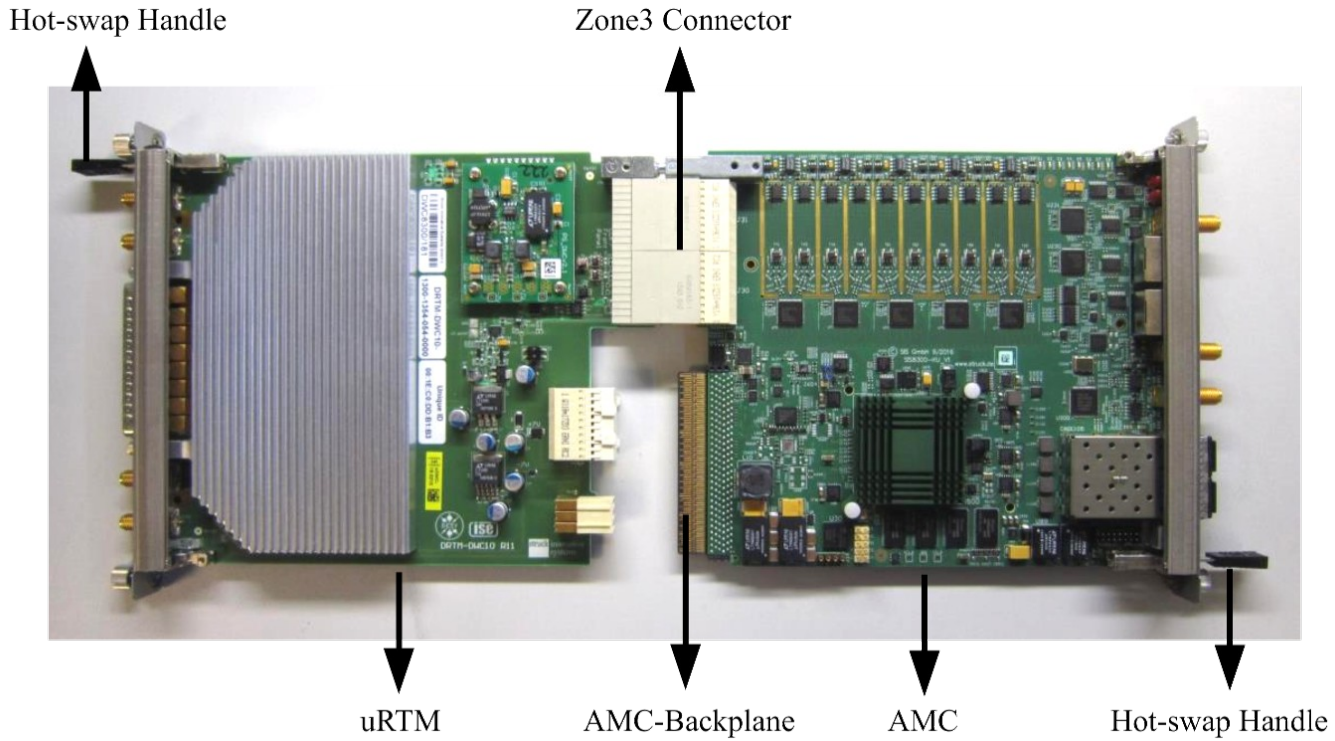

DWC8VM1 (RF Front End) + SIS8300L2 (ADC)

Power Limit: AMC + RTM < 80W

AMC Examples

- A/D & D/A Converters
- CPU
- FMC Carriers
- Timing Boards
- GPUs
- Storage (eg. SSD via SATA)
- Network Interface
- ...

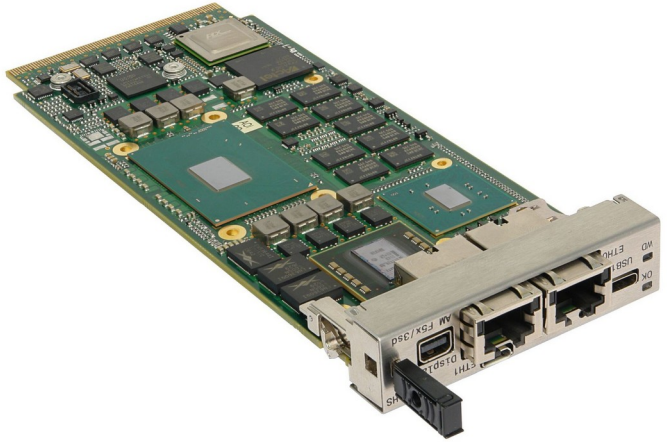

Concurrent AMC-CPU

DAMC-FMC2ZUP from MicroTCA Technology Lab

MCH

MicroTCA Carrier Hub

- Main **authority** of the crate
- Bare minimum functionality;
  - **Switch Fabric**
  - **Remote Control and Management Logic**
- Additonal functionalities;
  - Provide clocking functionality
  - Fat Pipe Switch (PCIe)
  - JTAG Connectivity to the backplane
- Can be ‘embedded’ inside the backplane. (Seen on smaller crates)
- Zone 3 Connectivity exists. (eg. MCH-RTM used as CPU)

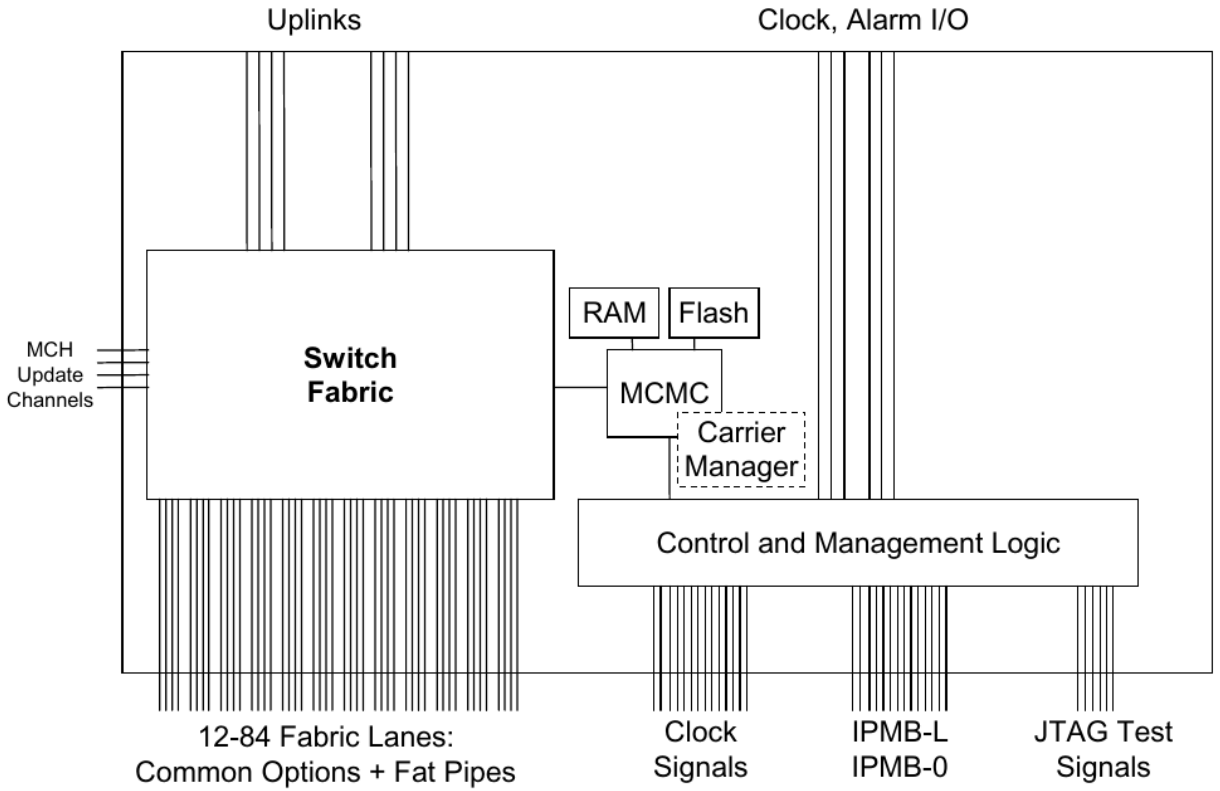

MicroTCA Carrier Hub

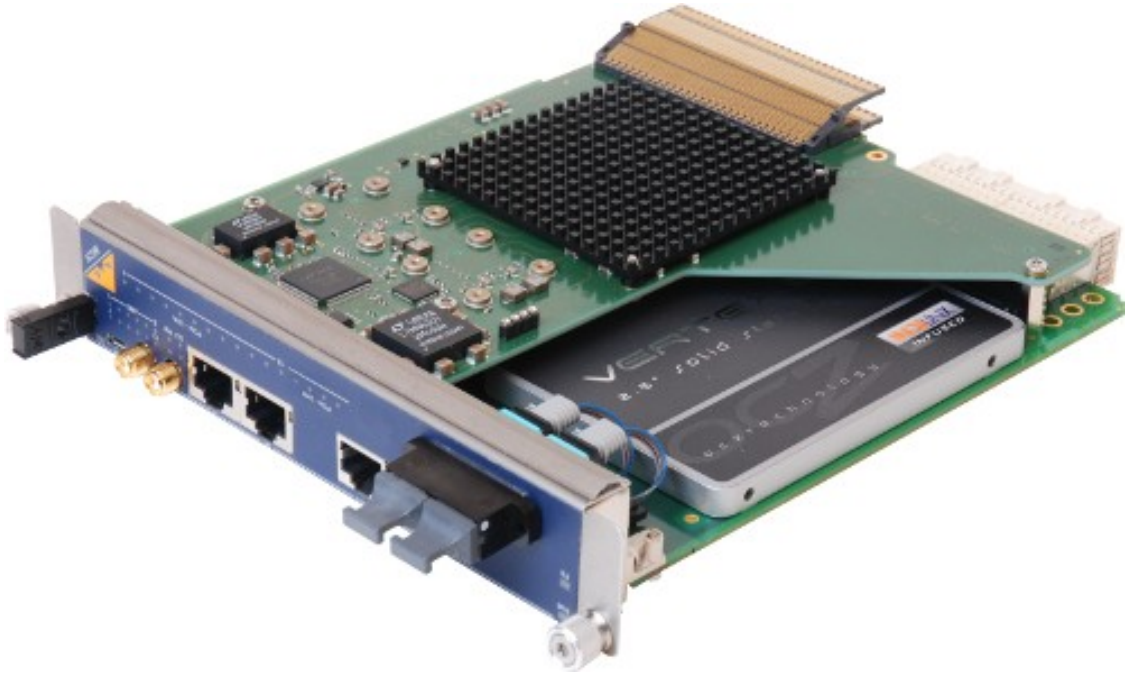

NAT-MCH-PHSY80 from N.A.T

UTC006 from Vadatech

Putting Everything together

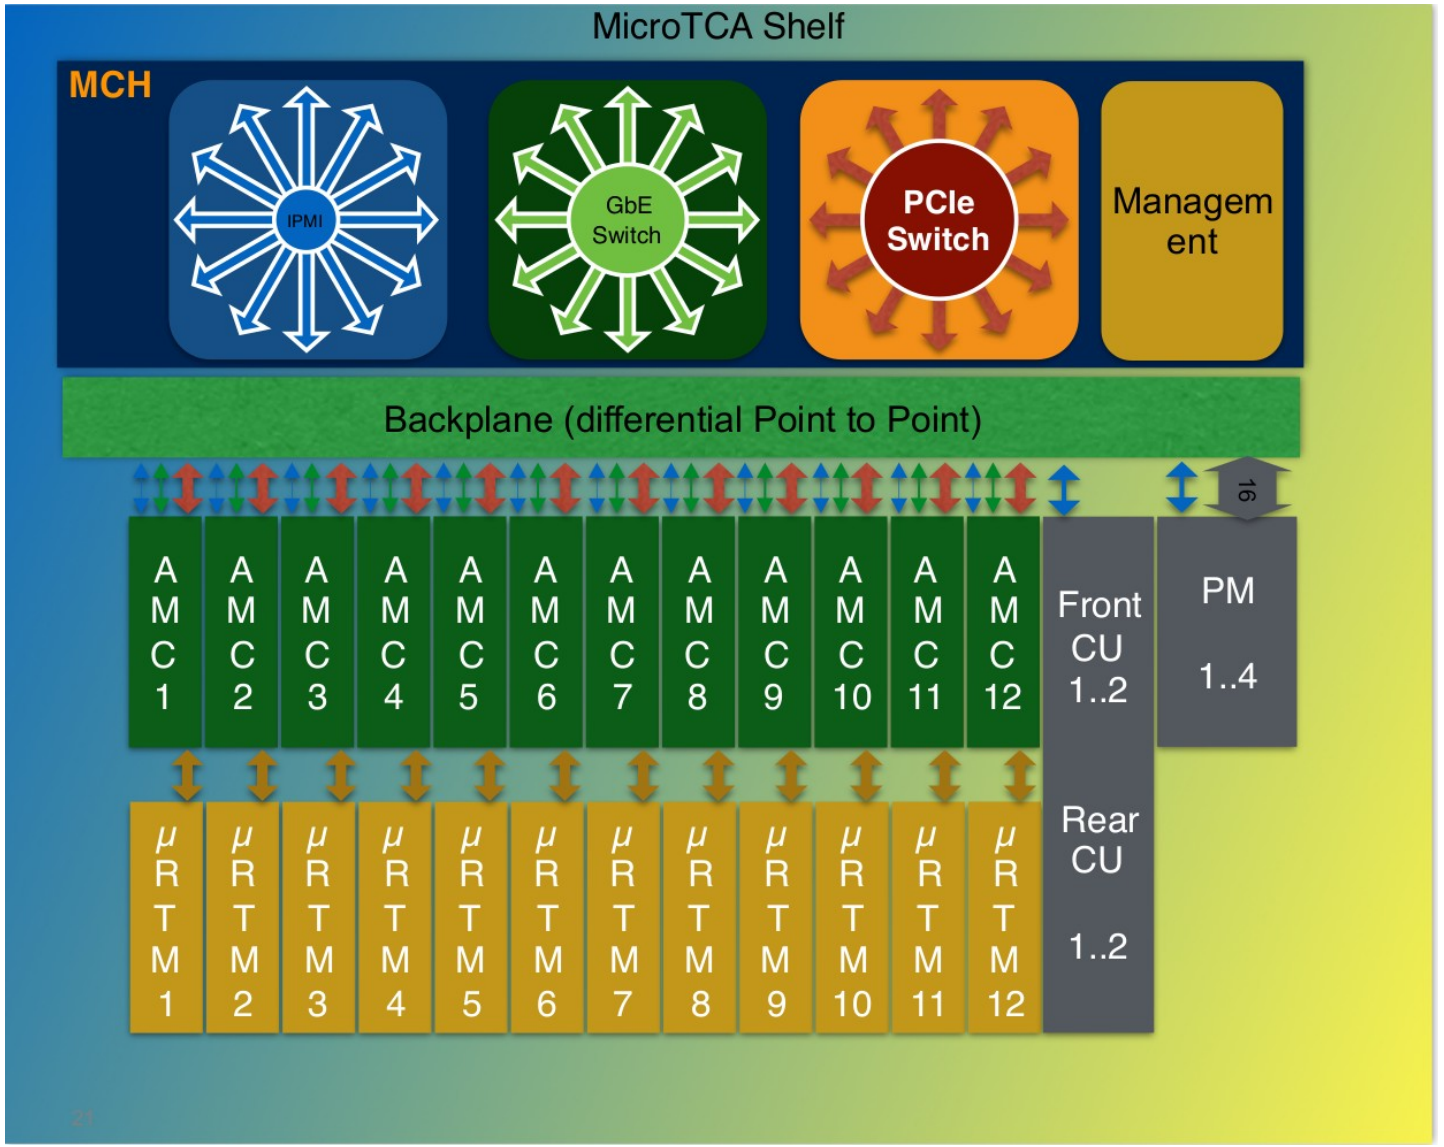

AMC-Backplane

The AMC-Backplane

- 170 pins grouped into **ports**
- Each port composed of 6 wires;
  - Differential TX
  - Differential RX
  - 2 GND Pins
- Protocols that can be seen on AMC Backplane:
  - IPMI (Management)
  - Gigabit Ethernet
  - SATA
  - FAT Pipe (PCIe |SRIO | 10/40GigE)
  - Point-to-Point Links
  - MLVDS Lanes
  - Clocks
  - JTAG

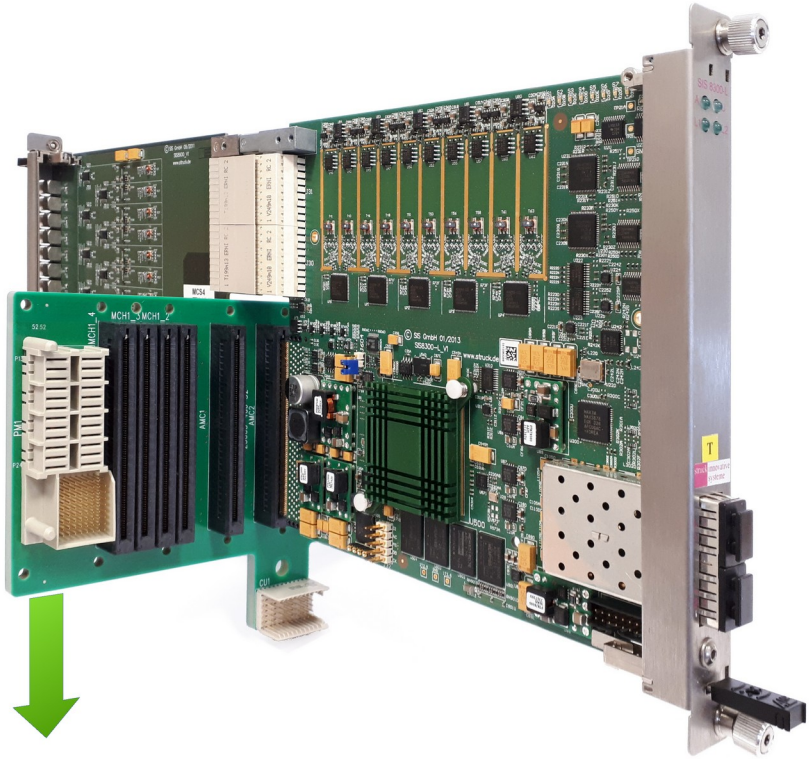

‘The Naked View’

RF Backplane (MicroTCA.4.1)

RF Backplane (MicroTCA.4.1)

Motivation: Getting rid of spaghetti, better management for analog signal distribution

Before

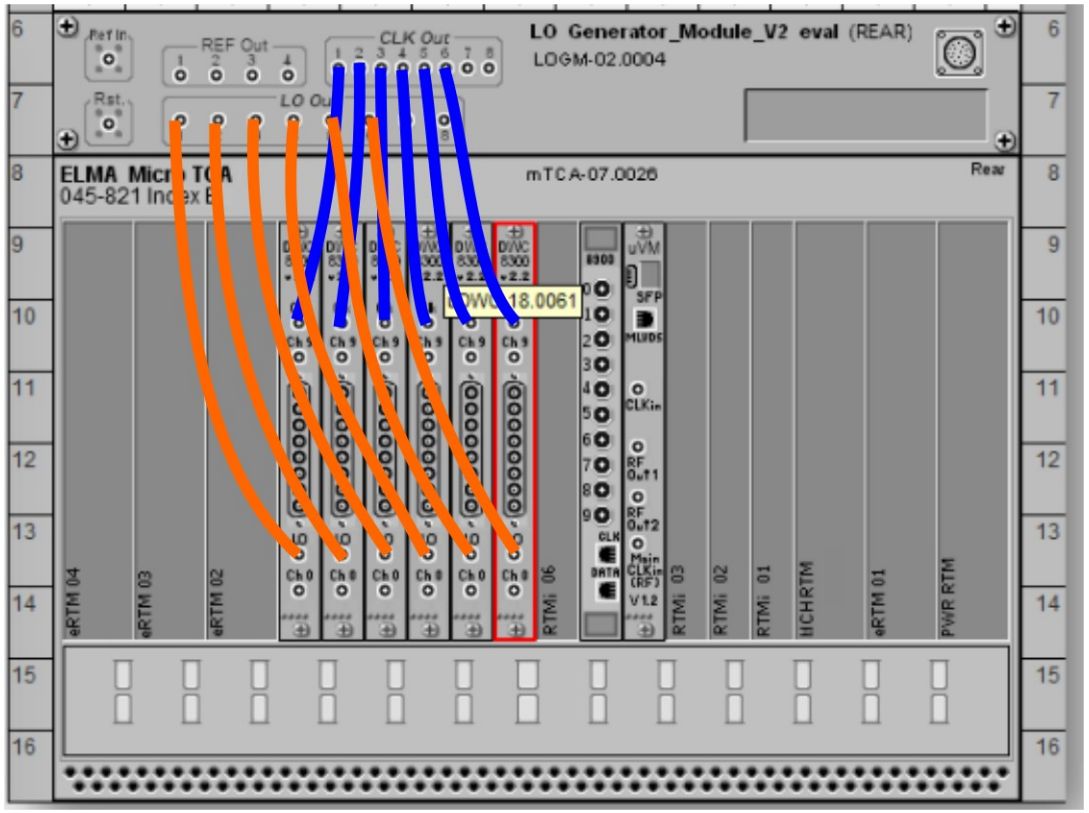

After

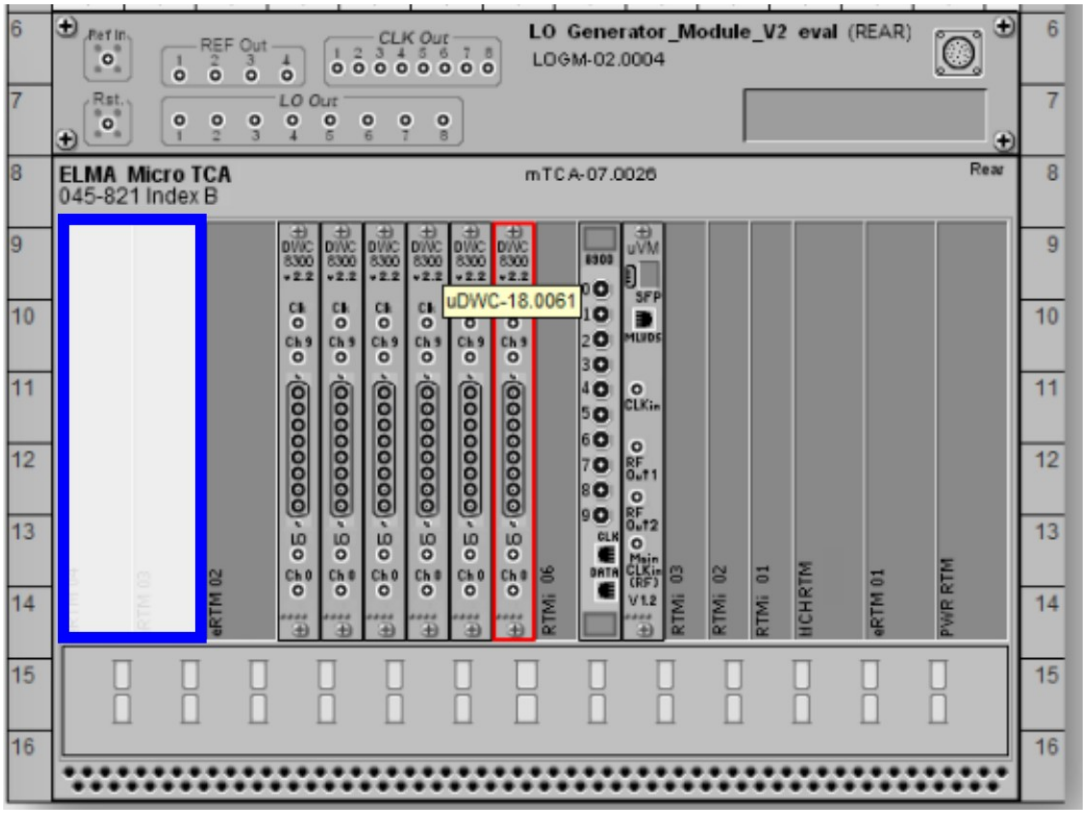

RF Backplane

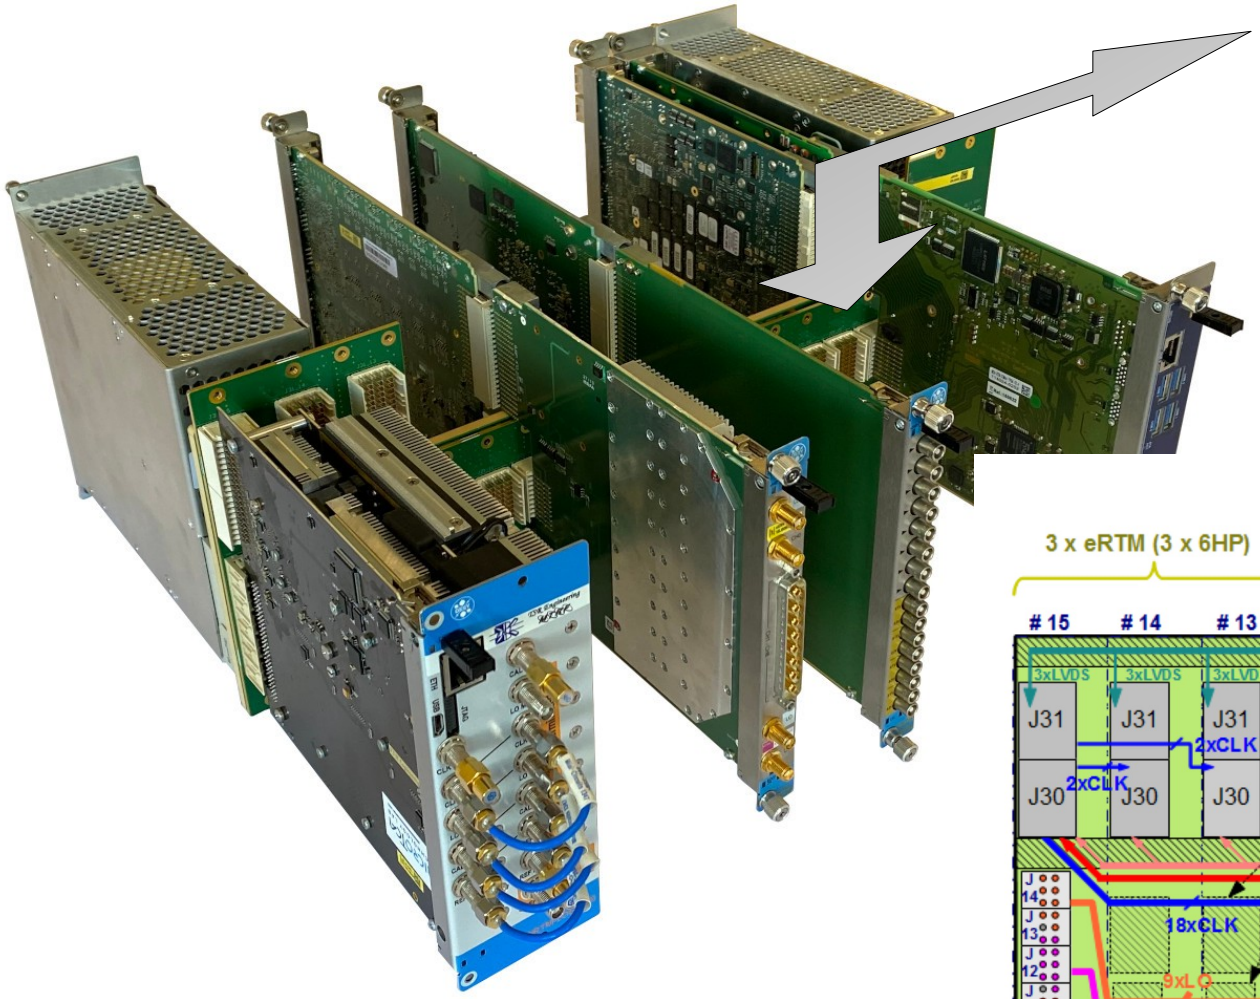

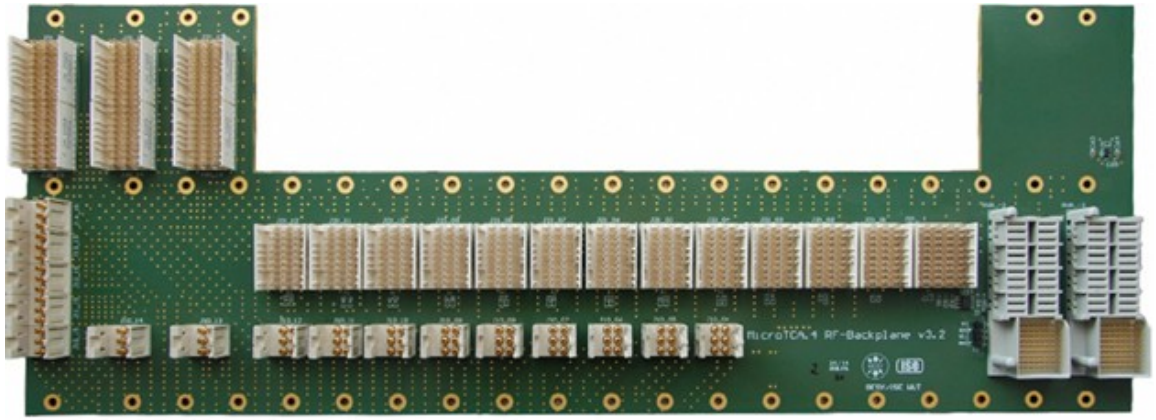

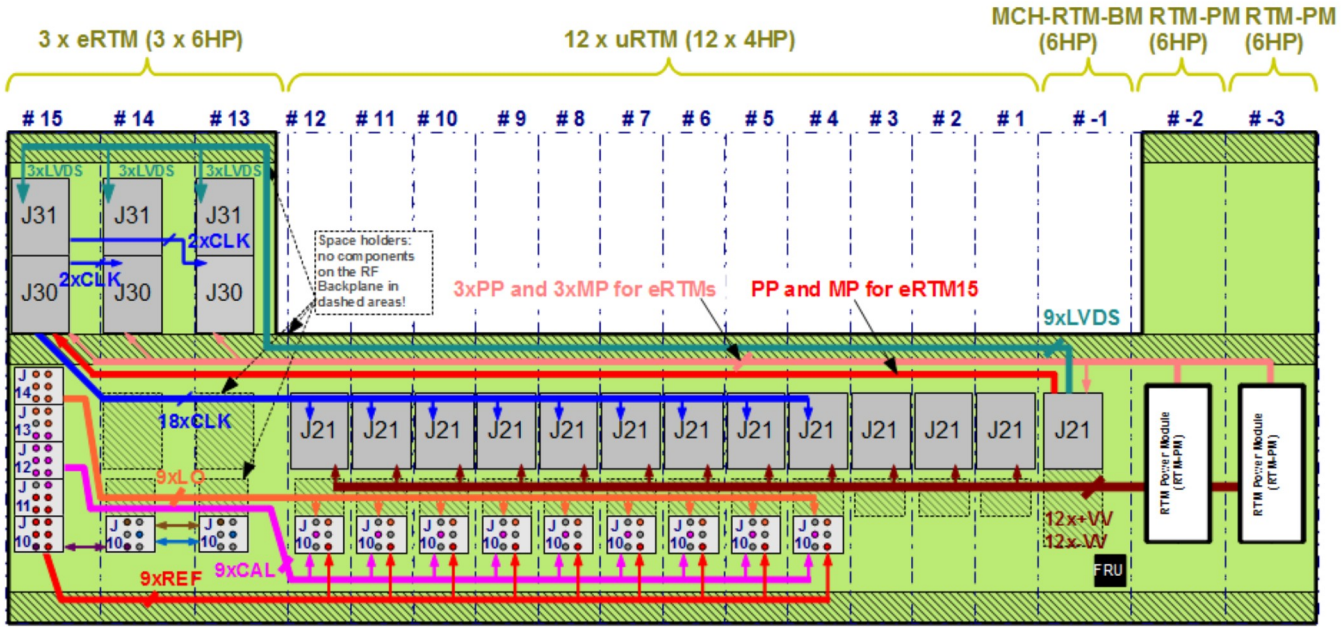

Management

Management

- **I**ntelligent **P**latform **M**anagement **I**nterface (**IPMI**) is used for management inside MicroTCA system
- Developed by Intel. Still used in servers for out-of-band management
- I2c based messages between master ↔ Slave
- MCH talks with other cards using IPMB lanes
- Shutdown/Start individual cards
- Gather sensor readings (Temperature,voltage/current readings etc.)
- Can be used for firmware upgrade

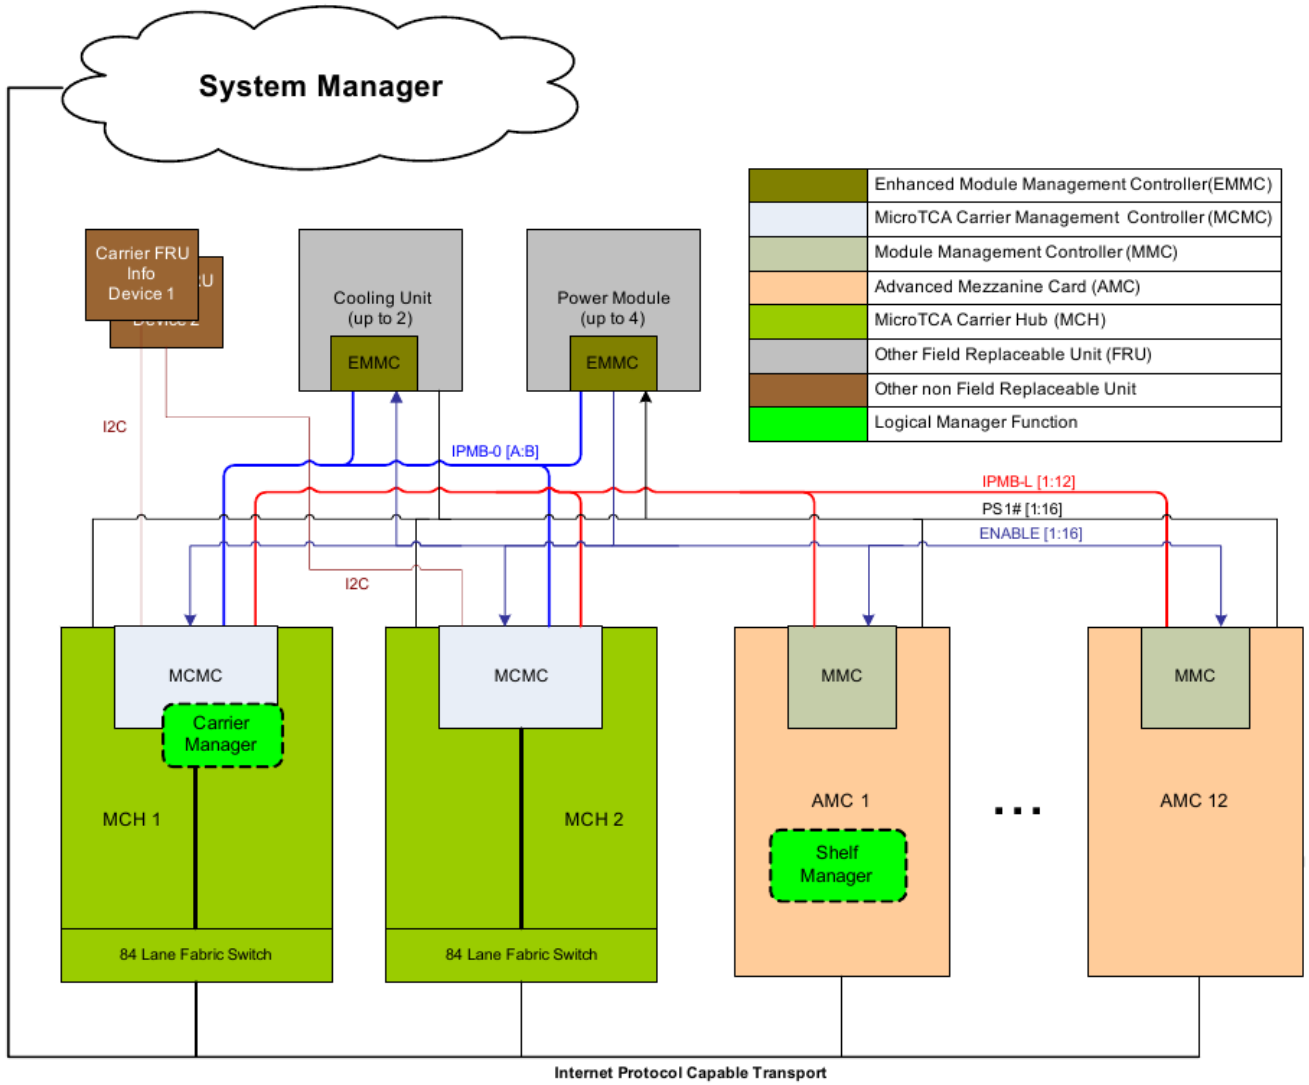

Management

- Controller: **MCMC**
  - **M**icroTCA **C**arrier **M**anagement **C**ontroller
- Agents: **MMC**
  - **M**icroTCA **M**anagement **C**ontroller
- Talking each other with I2c writes
- **Every AMC must have MMC implemented inside.**

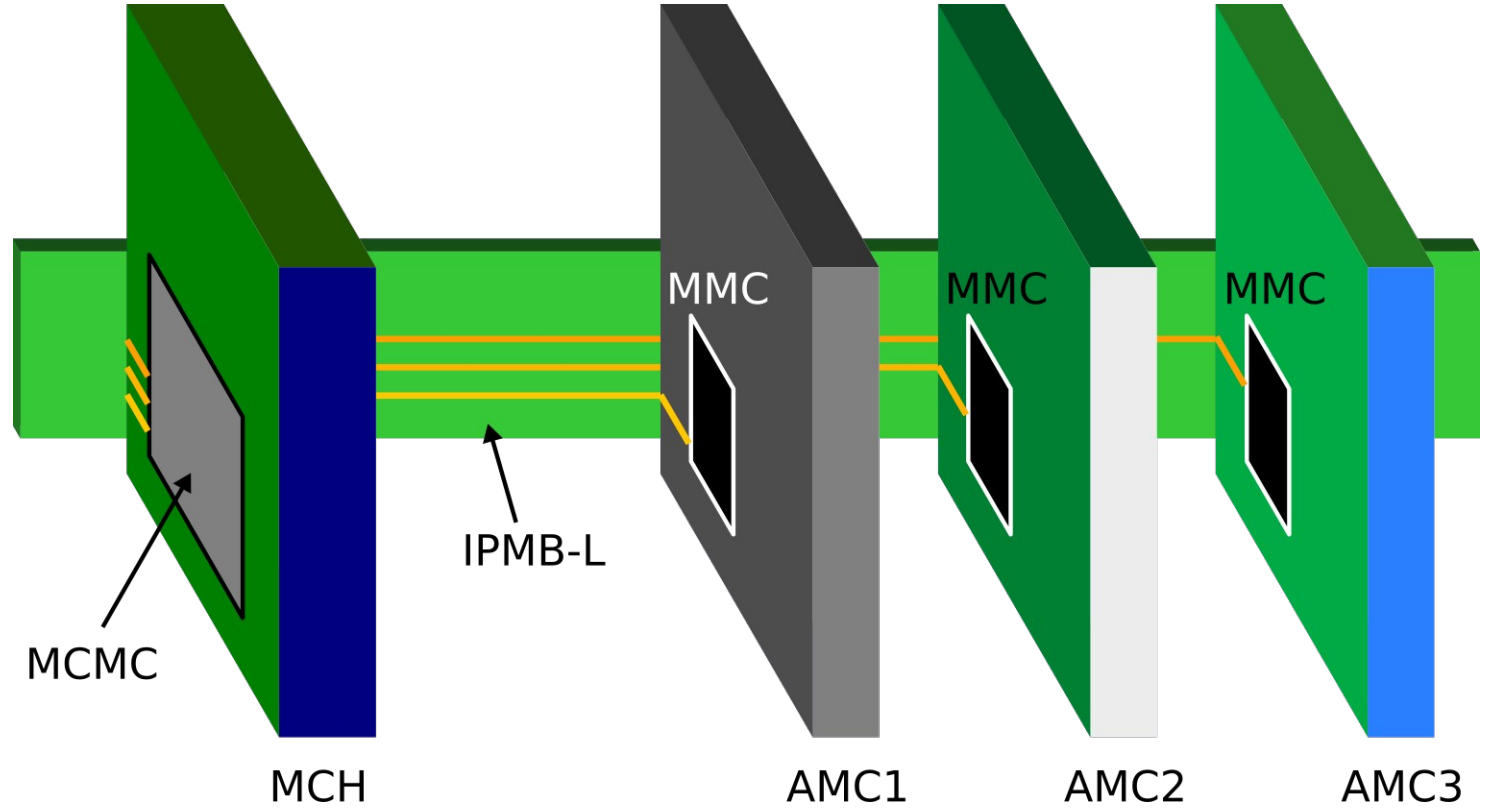

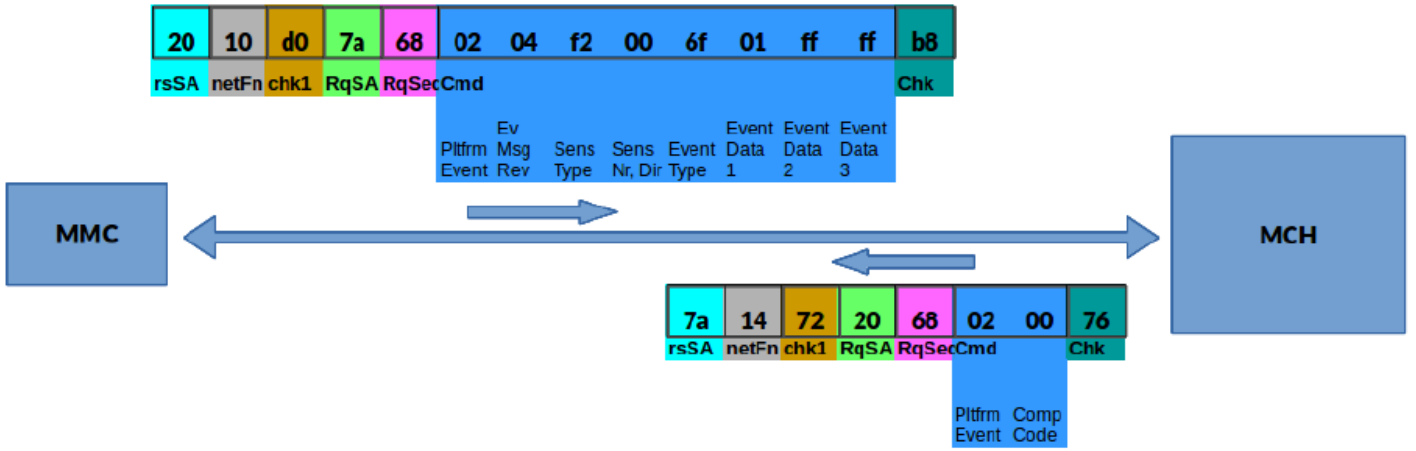

Field Replaceble Unit (FRU)

- The **passport** of the AMC
- Implemented as EEPROM connected to the MMC of the AMC
- MMC uses FRU to identify itself to the MCH
- Contains ‘fields’ of information such as:
  - Manufacturer
  - Board Information
  - Serial Number etc
- The AMC Backplane also has it’s own FRU.

```
nat> show_fruinfo 5
---------------------------------------
FRU Info for device 5:
---------------------------------------
Common Header    : 0x01 0x00 0x00 0x01 0x00 0x0a 0x00 0xf4
---------------------------------------
Internal Use Area : -
---------------------------------------
Chassis Info Area : -
---------------------------------------
Board Info Area          : at offs=8, len=72
Manufacturer(23)         : Concurrent Technologies
Board Name(10)           : AM 900/412
Serial Number(10)        : M23019/001
Part Number(11)          : 781-6011-13
FRU file ID(00)          :  -
---------------------------------------
Product Info Area : -
---------------------------------------
Multi Record Area  : at offs=80

Record(0): Type ID=0xc0, PICMG Record ID=0x16, offset=0x000, len=11
Module Current Requirements Record:
    Current Draw: 6.6 A

Record(1): Type ID=0xc0, PICMG Record ID=0x19, offset=0x00b, len=69
AMC Point-to-Point record:
AMC Slot  1, OEM GUID Count = 1
    Record Type = AMC, len=69
    Channel Descriptor count = 5
    Channel(0): Port[0 - - -]
    Channel(1): Port[1 - - -]
    Channel(2): Port[2 - - -]
    Channel(3): Port[3 - - -]
    Channel(4): Port[4 5 6 7]
    Link Descriptors: size=25
    Link 0 of Channel 0: lanes[0..3]=[1000], Eth ,  1000Base-BX, Grp=0x0, Match=0x0
    Link 1 of Channel 1: lanes[0..3]=[1000], Eth ,  1000Base-BX, Grp=0x0, Match=0x0
    Link 2 of Channel 2: lanes[0..3]=[1000], SATA,  LinkTExt=[1000], Grp=0x0, Match=0x2
    Link 3 of Channel 3: lanes[0..3]=[1000], SATA,  LinkTExt=[1000], Grp=0x0, Match=0x2
    Link 4 of Channel 4: lanes[0..3]=[1111], PCIe,  Gen 2, no SSC, Grp=0x0, Match=0x2

Record(2): Type ID=0xc0, PICMG Record ID=0x30, offset=0x050, len=15
Zone 3 Interface Compatibility record:
MTCA.4 REP Number: 0x44332211
---------------------------------------
```

Example of a FRU (CPU-AMC)

AMC-Backplane

Protocols on the AMC-Backplane

- IPMI (Management)
- Gigabit Ethernet (Ports 0-1)
- SATA (Ports 2-3)
- Fat Pipe + Extended Fat Pipe (Ports 4-11)
  - PCIe
  - SRIO
  - 10/40 GbE
- Point-to-Point Links (Ports 12-15)
- MVLDS (Ports 4-11)
- Clocks (TCLKA,TCLKB TCLKC,FCLK)
- JTAG

Figure 6-11 AMC Port mapping regions

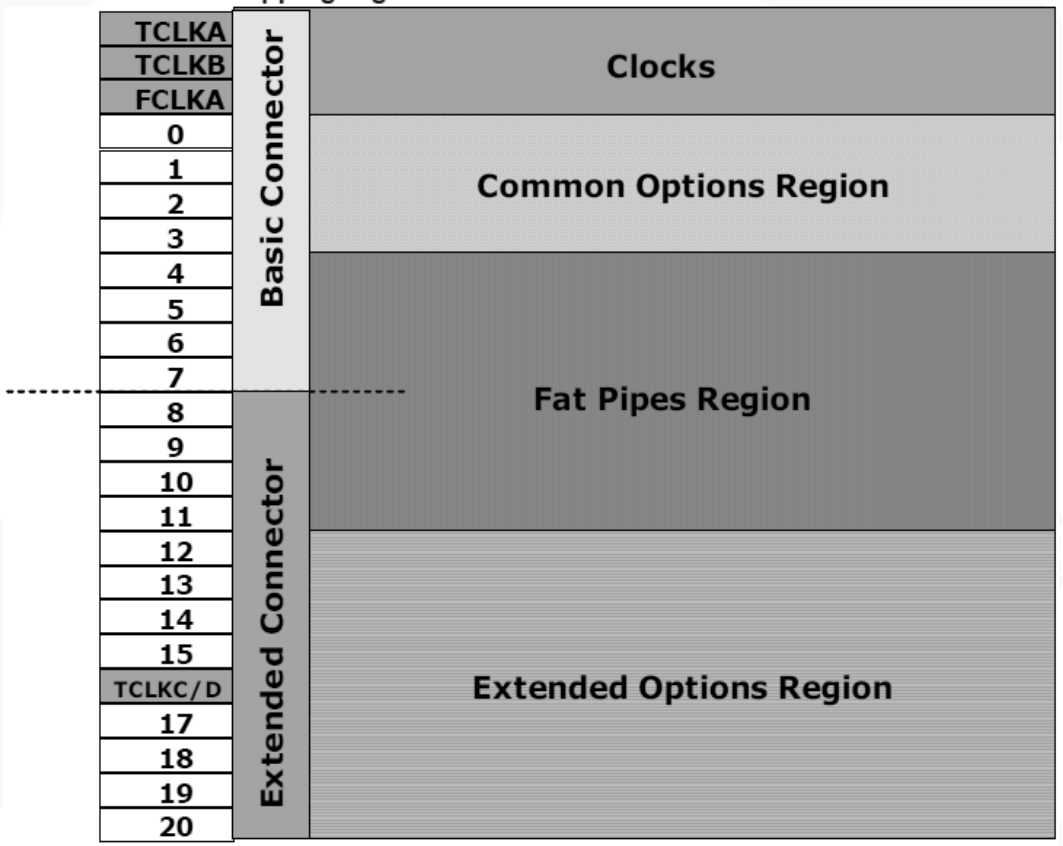

Protocols on the AMC-Backplane

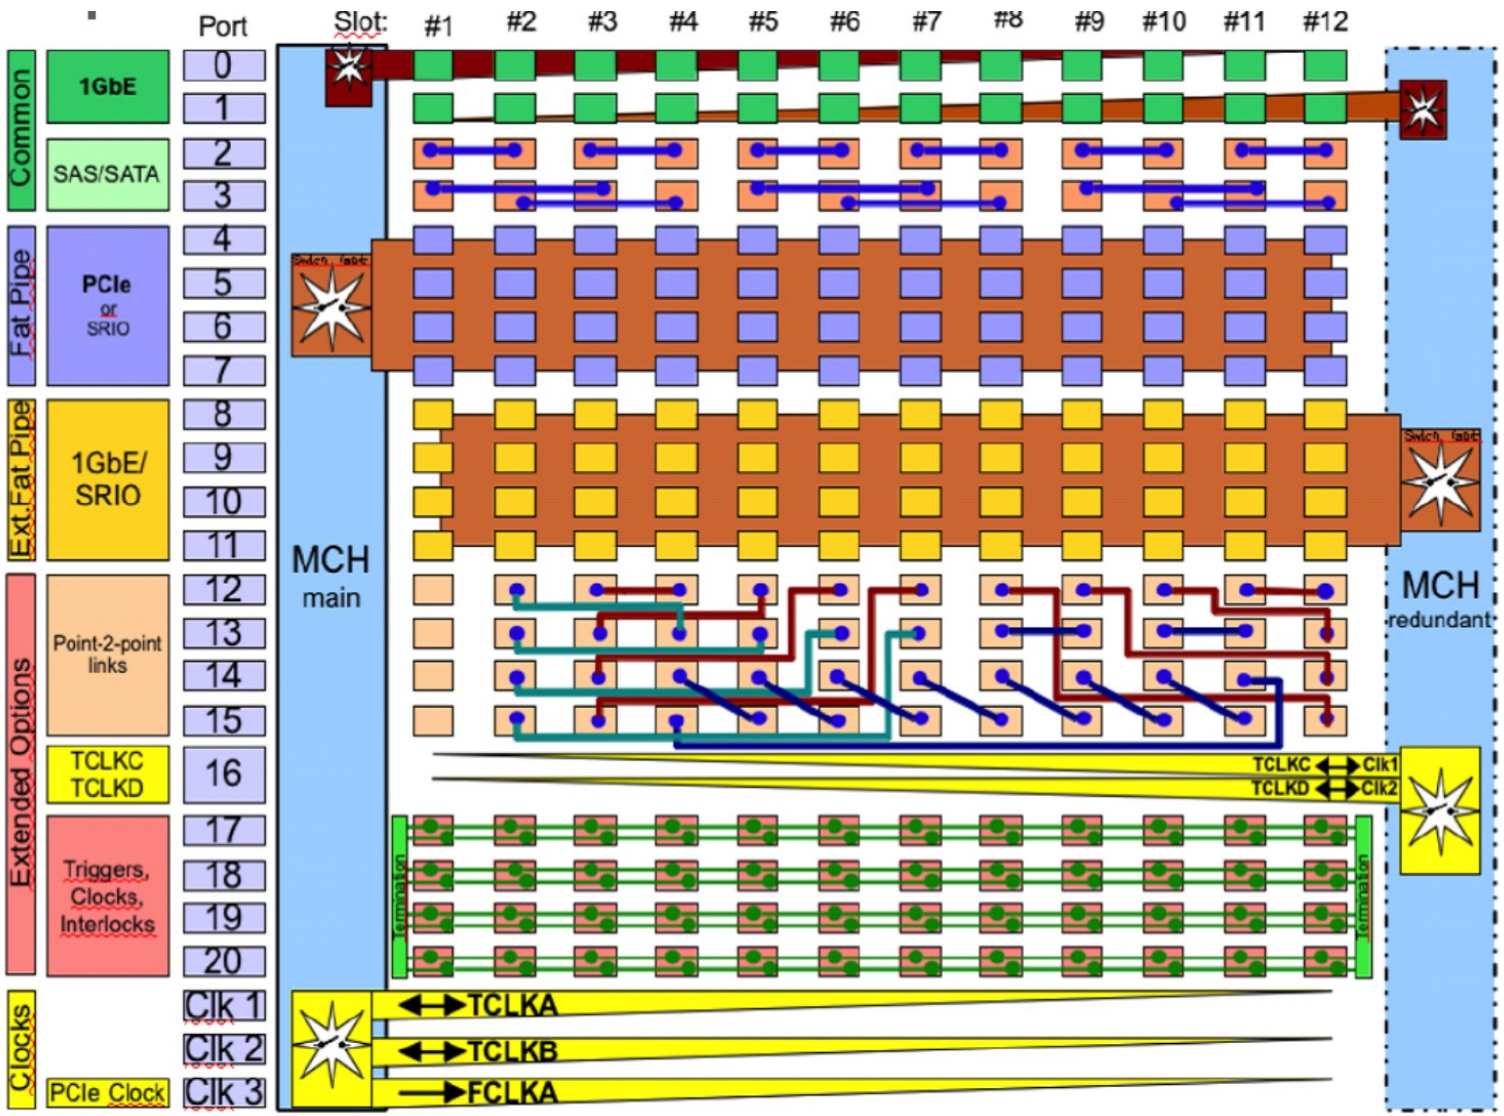

Point-to-Point Links

8 x M-LVDS: Trigger

PCIe

PCIe

- MicroTCA standard offers PCIe lanes **up-to x8 (Gen 3)** connectivity between cards
- MCH holds the PCIe switch and can talk the switch
  - Choose any-slot to be the ‘root-complex’. (eg. CPU)
  - See the link status online
  - Configure the PCIe switch parameters
  - Multiple root-complex inside single crate possible
- **Hot Swap is possible!** (Take the AMC out, re-insert again and things should continue work)

PCI Express link performance[46][47]

Version	Intro-duced	Line code	Transfer rate[i][ii]	Throughput[i][iii]				
				×1	×2	×4	×8	×16
1.0	2003	8b/10b	2.5 GT/s	0.250 GB/s	0.500 GB/s	1.000 GB/s	2.000 GB/s	4.000 GB/s
2.0	2007	8b/10b	5.0 GT/s	0.500 GB/s	1.000 GB/s	2.000 GB/s	4.000 GB/s	8.000 GB/s
Now 3.0	2010	128b/130b	8.0 GT/s	0.985 GB/s	1.969 GB/s	3.938 GB/s	7.877 GB/s	15.754 GB/s
4.0	2017	128b/130b	16.0 GT/s	1.969 GB/s	3.938 GB/s	7.877 GB/s	15.754 GB/s	31.508 GB/s
Future 5.0	2019	128b/130b	32.0 GT/s	3.938 GB/s	7.877 GB/s	15.754 GB/s	31.508 GB/s	63.015 GB/s
6.0 (planned)	2021	128b/130b + PAM-4 + ECC	64.0 GT/s	7.877 GB/s	15.754 GB/s	31.508 GB/s	63.015 GB/s	126.031 GB/s

PCIe

How the data moves inside the crate using PCIe protocol?

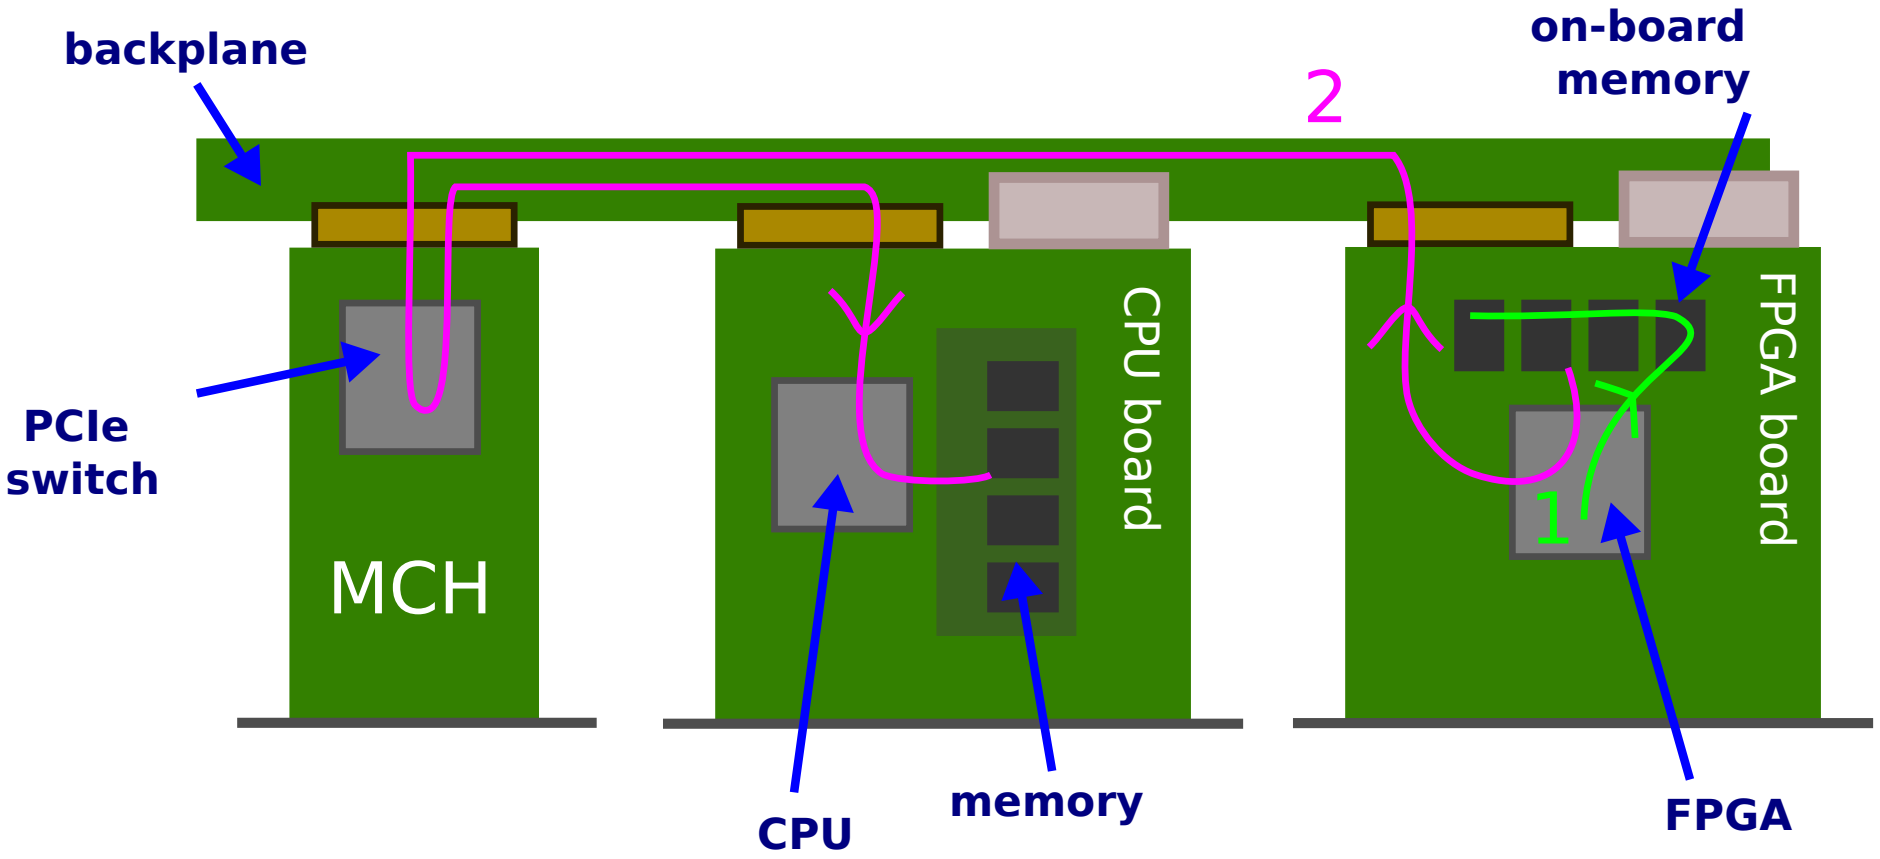

Protocols on AMC-Backplane

Point-to-Point Links

Point-to-Point Links

Point to Point links offer direct communication from FPGA to FPGA.

Used for data aggregation / fast feedback between boards

These lines are ‘hard wired’. Double check the connectivity before ordering.

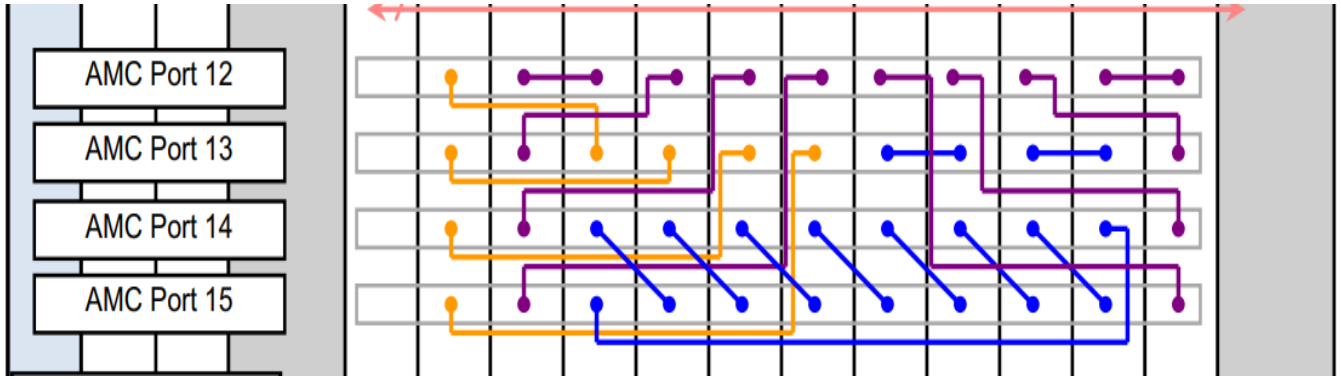

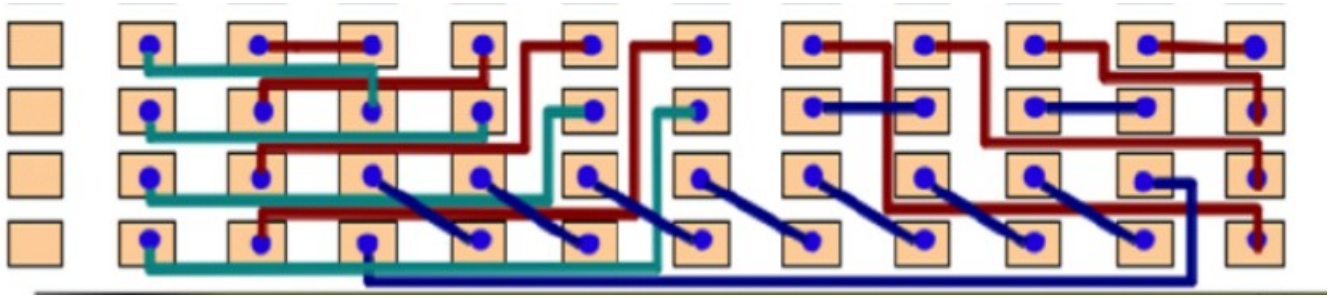

Point-to-Point Links

Example from European-XFEL

- Data aggregation on point-to-point links on Europan-XFEL LLRF Crates:
- Probe + Forward + Reflected signals of 16 cavities gets send to main controller board.
- Some numbers:
  - 6.25Gbps link rate
  - Sending 11x32 bits payload packet
  - End to End latency: ~344ns

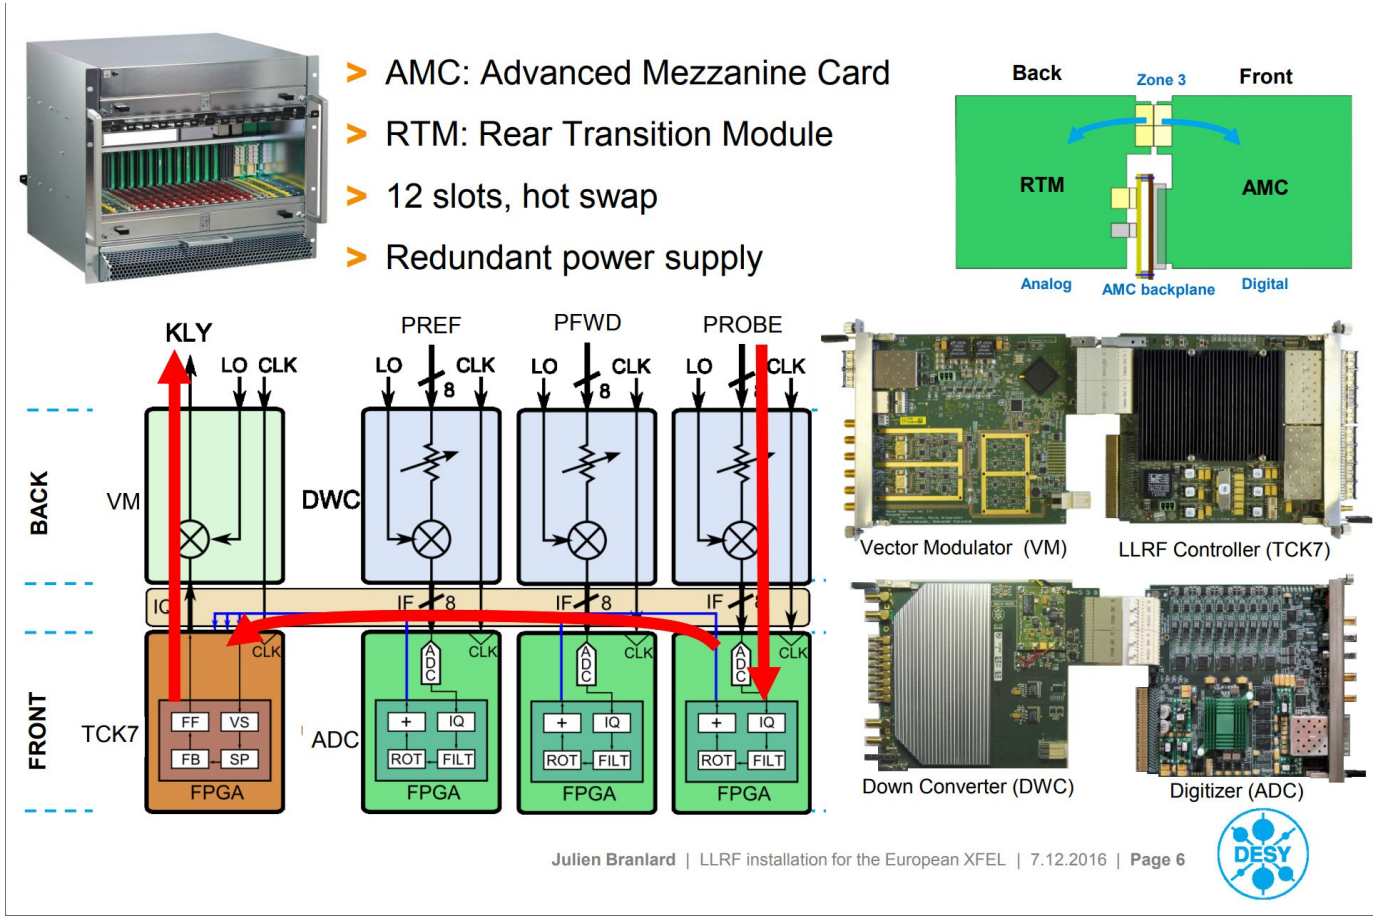

MLVDS

MLVDS

- Ports 17-20 Can be used to forward clocks, triggers and interlock to all other cards on the crate.
- Mesh Topology
- Multipoint LVDS is used in MicroTCA for communication between cards. On each individual line one card (application specific) acts as a driver, other cards can be configured as receivers.
- Wired OR is also possible in MLVDS - more than one card can drive the same line (with the same polarity)

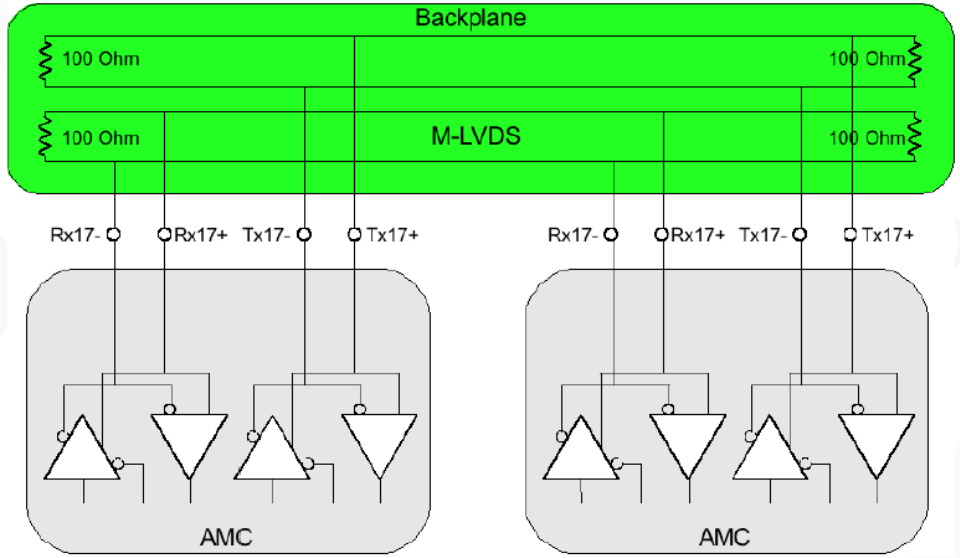

Figure 6-4: M-LVDS transceiver shown for port 17

Table 6-1: Example usage of the 8 bus lines for triggers, interlocks and clocks

AMC Port	Name	Description	Usage
Rx17	TrigStart	Start sampling data	Triggers
Tx17	TrigEnd	Stop sampling data	
Rx18	TrigReadOut	Start data transfer to CPU	
Tx18	ClkAux	Low performance clock	
Rx19	Reset	Reset of counter, dividers	
Tx19	Interlock 0	Interlock line 0	3 interlocks to provide 2 out of 3 redundancy
Rx20	Interlock 1	Interlock line 1	
Tx20	Interlock 2	Interlock line 2	

Data Transfer inside Crate

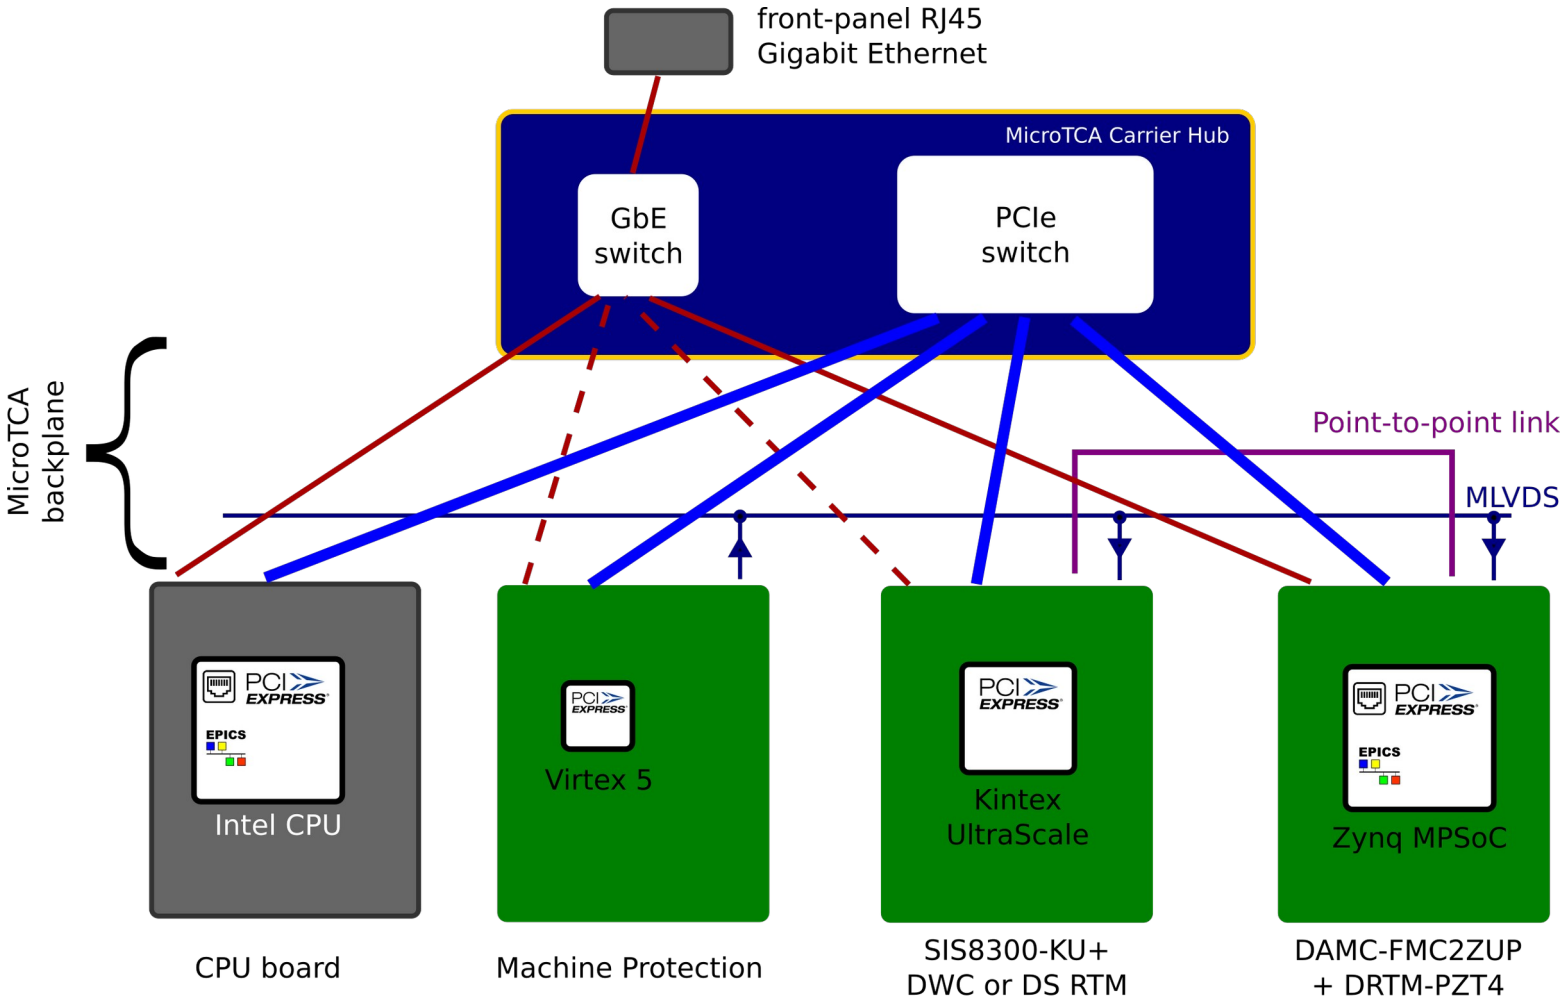

Do you want to learn more?

- Go to techlab.desy.de to learn more about the training
- 2 Trainings:
  - Basic
  - Advanced
- Dates for 2022 will be announced soon!
- Training can be held virtually/in-house depending on the health guidance.

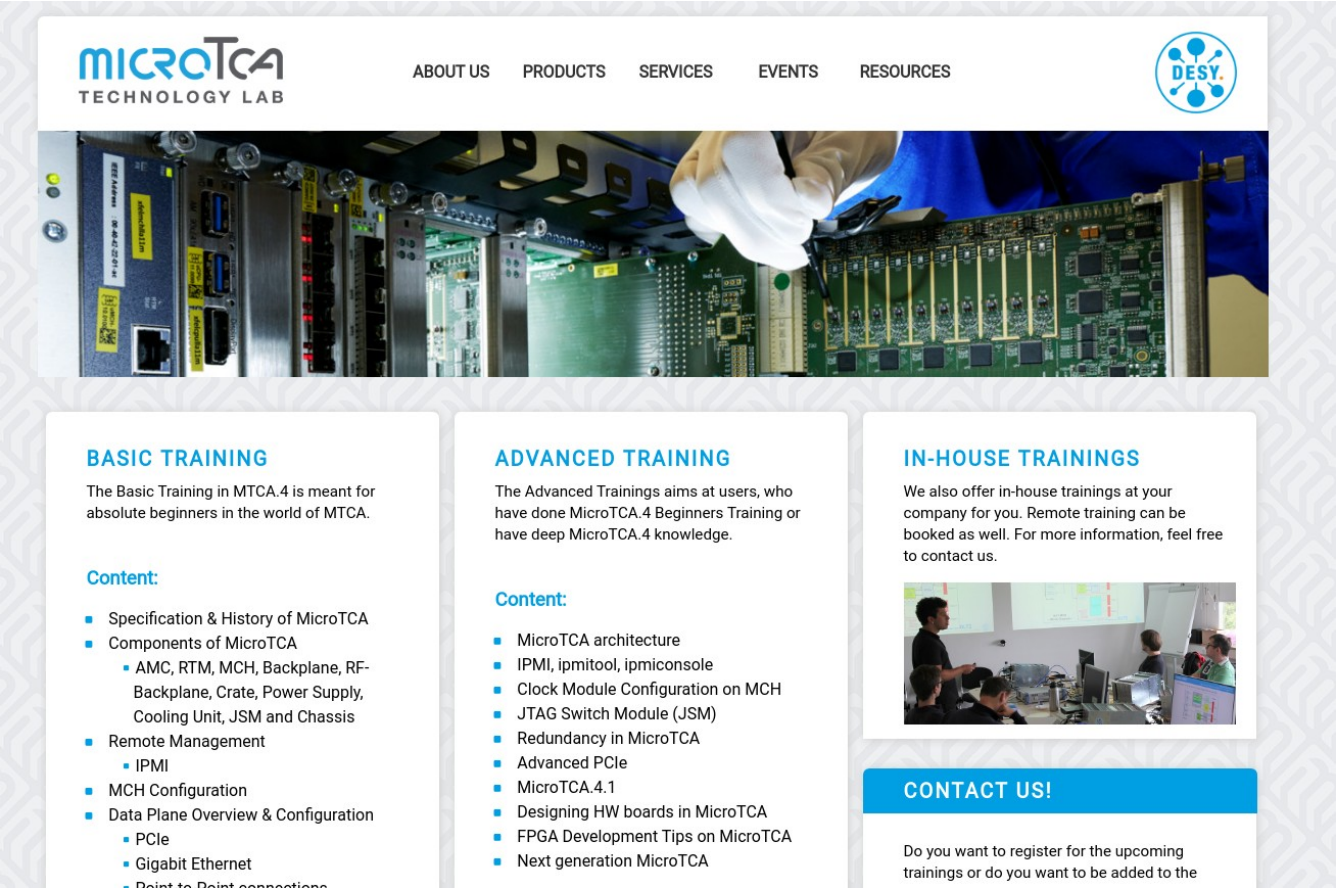

Thank you

Contact

DESY. Deutsches
Elektronen-Synchrotron

www.desy.de

Cagil Gumus (CJ)
MSK- Firmware Team
cagil.guemues@desy.de

Extra Things

MicroTCA Carrier Hub

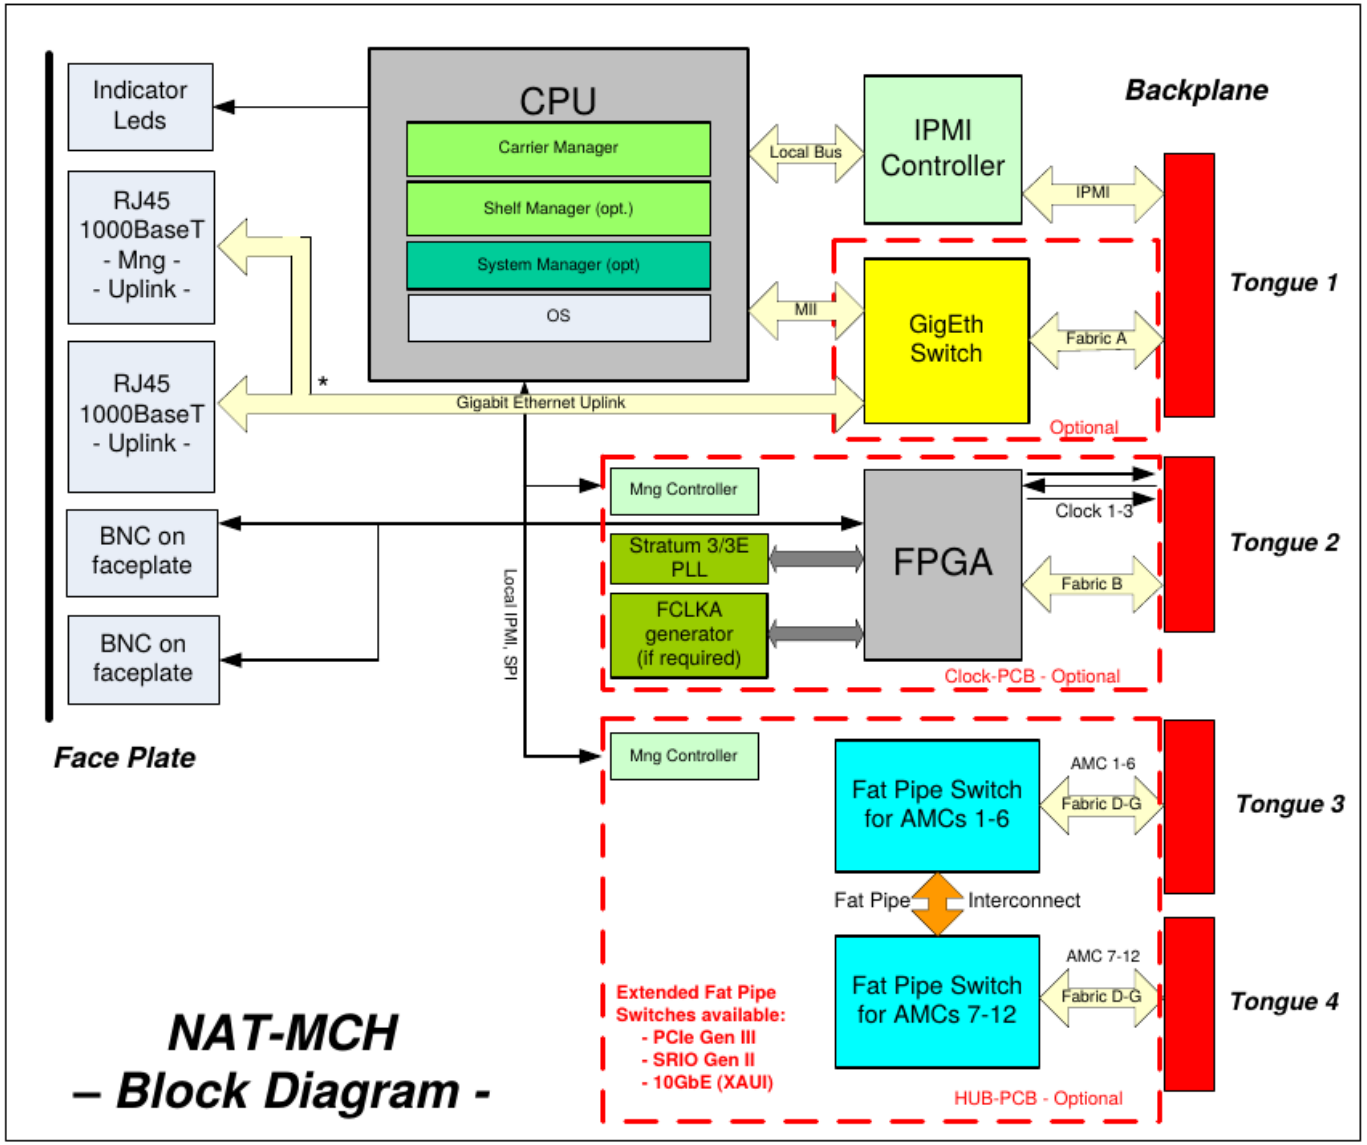

PCIe Root Complex outside of the crate

Suffering from weak CPU-AMC? Here is your solution

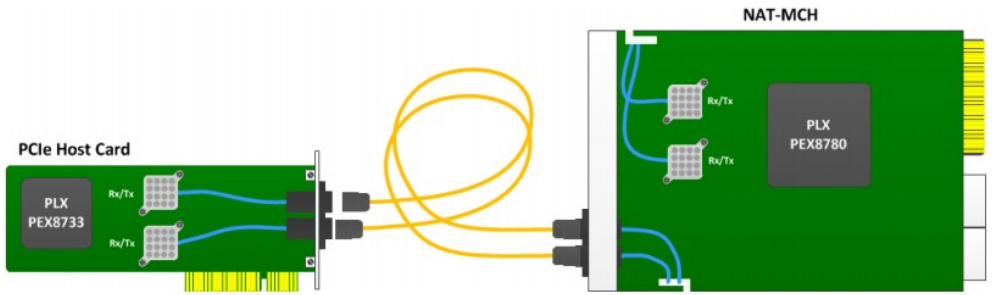

Needed Parts:

- 4 x Finisar BOA
- 4 x Pig Tail
- 4 x Face Plate Adapter
- 2 x Patch Cord 5m
- Resulting Costs for a PCIe

GenIII x16 Uplink Connection:

Pros:

- Cheaper & Poweful PC outside of 80W limitation
- Many choices in the industry for parts
- Many more PCIe slots available on the motherboard for more cards

Cons:

- CPU is not managed by MCH
- Boot sequence of crate and PC has to be done properly

PCIe Tree

‘lspci -tv’ command visualised