

MTCA.4 TUTORIAL BASICS INTRODUCTION IN XTCA

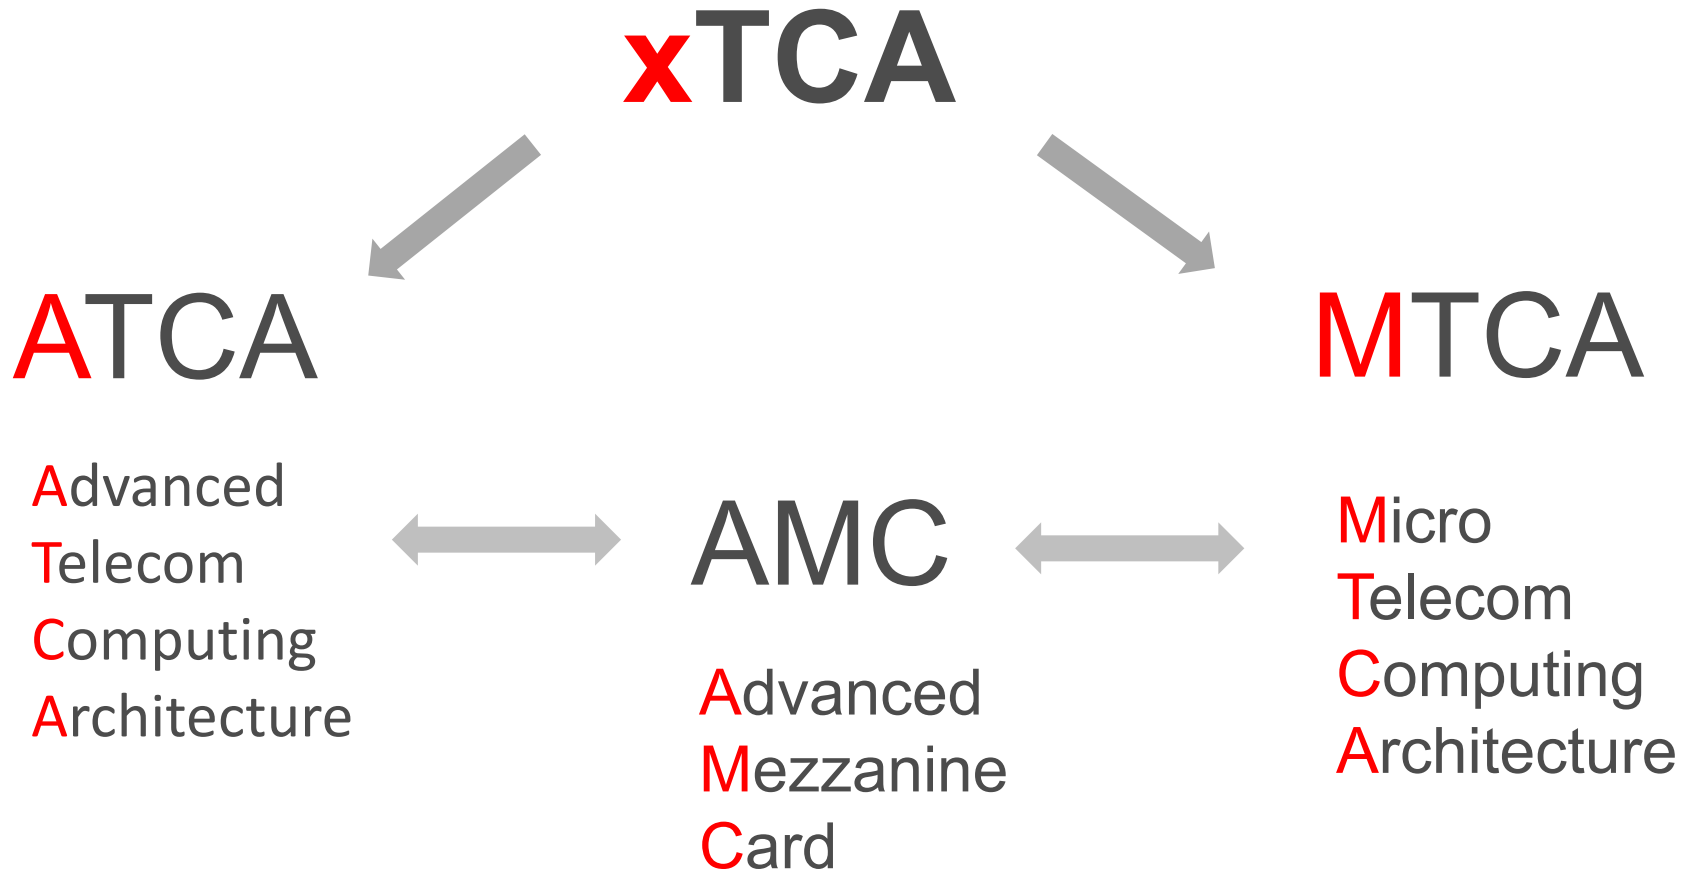
6TH MTCA WORKSHOP FOR INDUSTRY AND RESEARCH
DECEMBER, 2017
DESY, HAMBURG

Ralf Waldt
Pentair Technical Solutions GmbH
ralf.waldt@pentair.com

AGENDA

- What is xTCA?
- Specifications Overview
- ATCA Features
- AMC Features
- MTCA.0 Features
- MTCA.4
 - Initial Requirements
 - Mechanical Features
 - Module sizes
 - Management extensions compared to MTCA.0
 - Keying
 - Hot Swap Transition States
 - Backplane
 - Cooling
 - Redundancy

What is xTCA?

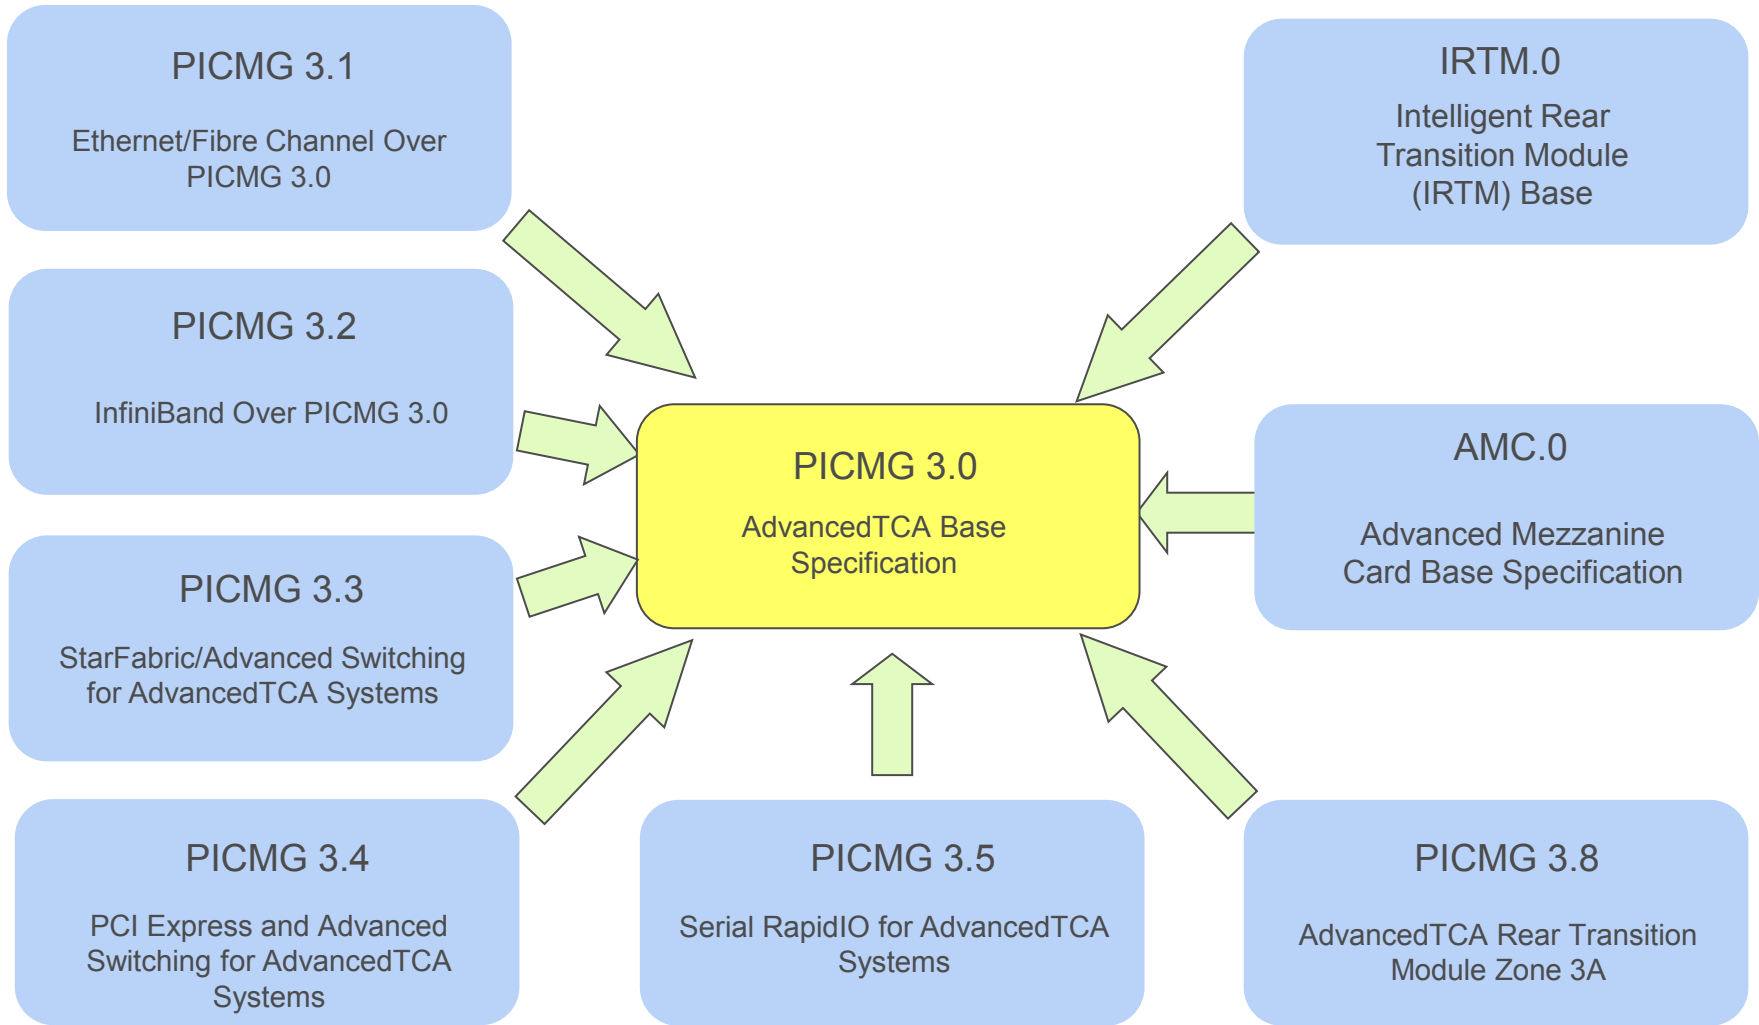


ATCA

Overview

- Specification initially targeted to the Telco Industry
- Industry, however, HA features of interest to
- Many other “up time” critical systems!
- Introduced in 2002
- System throughput to 2 Tb/s (full mesh)!
- System Availability 99.999% (~5 min/yr)!
- Port data rate to 40 Gb/s (4 x 10Gb/s), 100Gb/s is now available!
- Management, monitoring and control!
- Software infrastructure providing API's, etc!
- Sponsored by the PCI Industrial Computer
- Manufacturers Group (PICMG)


ATCA Specification Family



ATCA Crate Elements

Front View

Rear View



Alarm Board

Backplane

Redundant Fan Trays

Dual redundant Shelf Managers




Power Entry Module (PEM)

ATCA Crate Elements

Board size and connectors

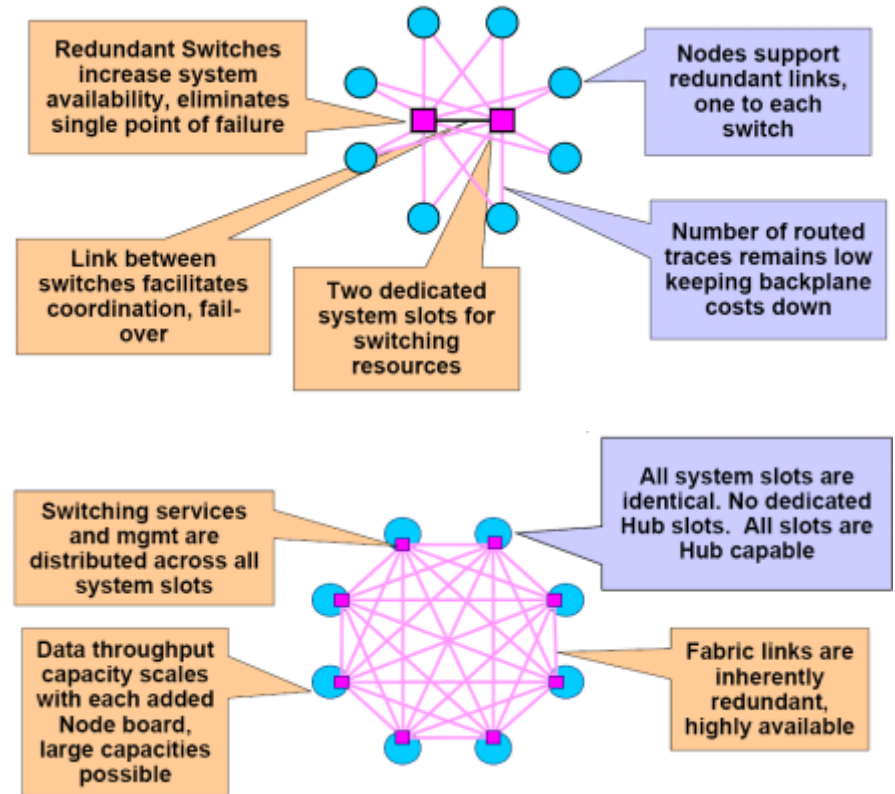
- Front board size 8U x 280
- Rear board (RTM) size 8U x 70 mm
 - Connects directly to front board
- Board width 6HP (1.2")
- Alignment/Key pins
- Zone 1: Management and Power
- Zone 2: Base Interface and Fabric Interface
- Zone 3: Interface to RTM



ATCA Interfaces

Zone 2 Backplane Interfaces

- Base Interface
 - 10/100/1000 BASE-T Ethernet
 - Always Dual Star topology
- Fabric Interface
 - Star topology
 - Mesh topology
- Clock Interface
 - Three dedicated clock interfaces
- Update Channel
 - Direct connection between two slots



ATCA Management

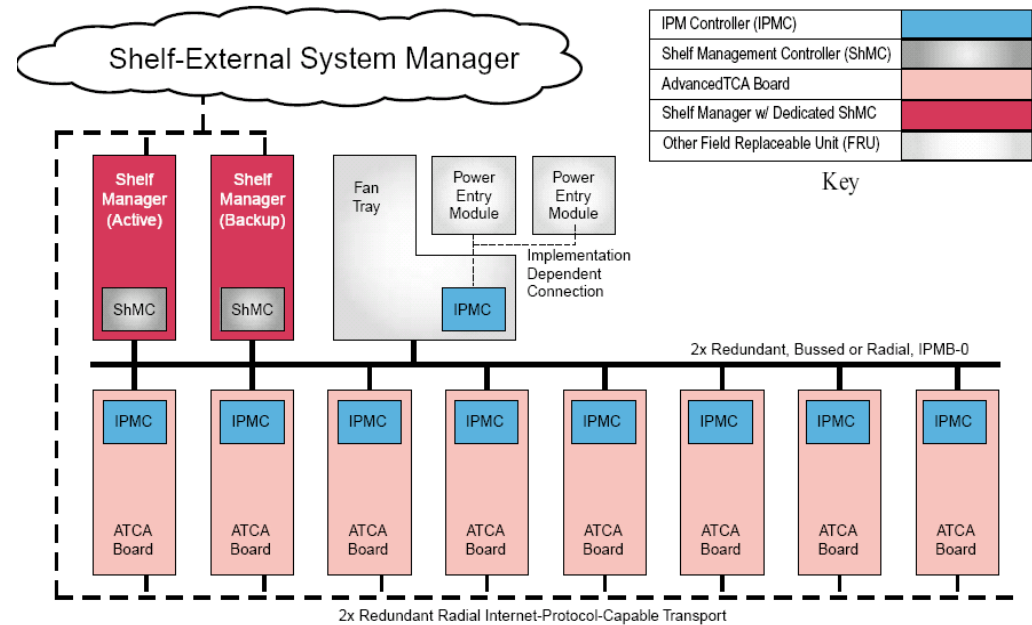
ATCA Shelf Management purpose

- Monitor & control low-level aspects of ATCA boards and other Field Replaceable Units within a shelf
- Watch over basic health of the shelf, report anomalies, take corrective action when needed
- Retrieve inventory information & sensor readings
- Receive event reports and failure notifications from boards and other intelligent FRUs
- Manage power, cooling & interconnect resources in the shelf (electronic keying)
- Management Protocol IPMI (I2C-bus on backplane)



ATCA Management

- Dedicated Shelf Management Controller (ShMC)
- ATCA Boards with IPMC
- Protocol IPMI (Physical layer I²C-Bus)
- Intelligent and Managed FRUs
- Bused or Radial IPMB




AMC


AMC Modules

- Initially developed as function extension for ATCA Boards
- Fully integrated into the ATCA IPMI management structure
- Plugged into a so called ATCA Carrier
- Hot Swap capability

AMC Modules






AMC Carrier



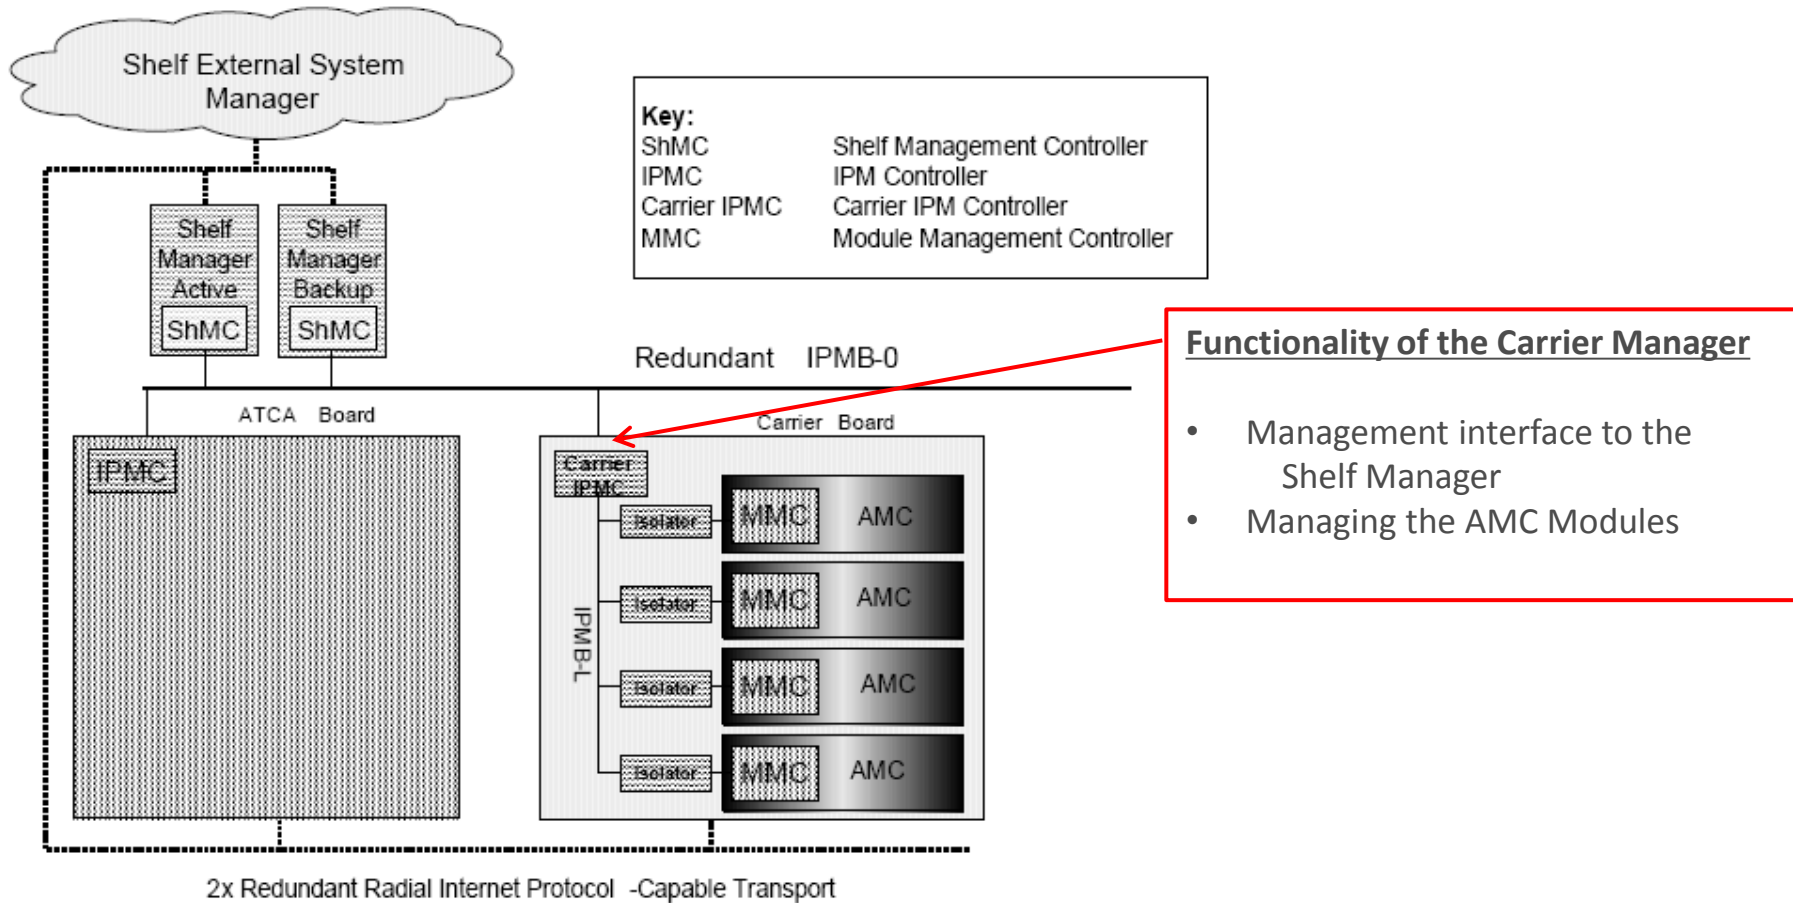
AMC Modules

AMC Module Sizes



AMC Modules


- Carrier IPMC represents the MMC on the AMC as a FRU to the Shelf Manager



MTCA.0

MTCA.0

- The basic idea of MTCA is to have a shelf that contains just AMC modules
- Backplane directly accepts AMC modules
- AMCs are interchangeable between ATCA and MTCA
- The infrastructure of a ATCA Carrier was adapted into the MTCA shelf (power, management, switching)
- No rear I/O, power input and all outputs to the front

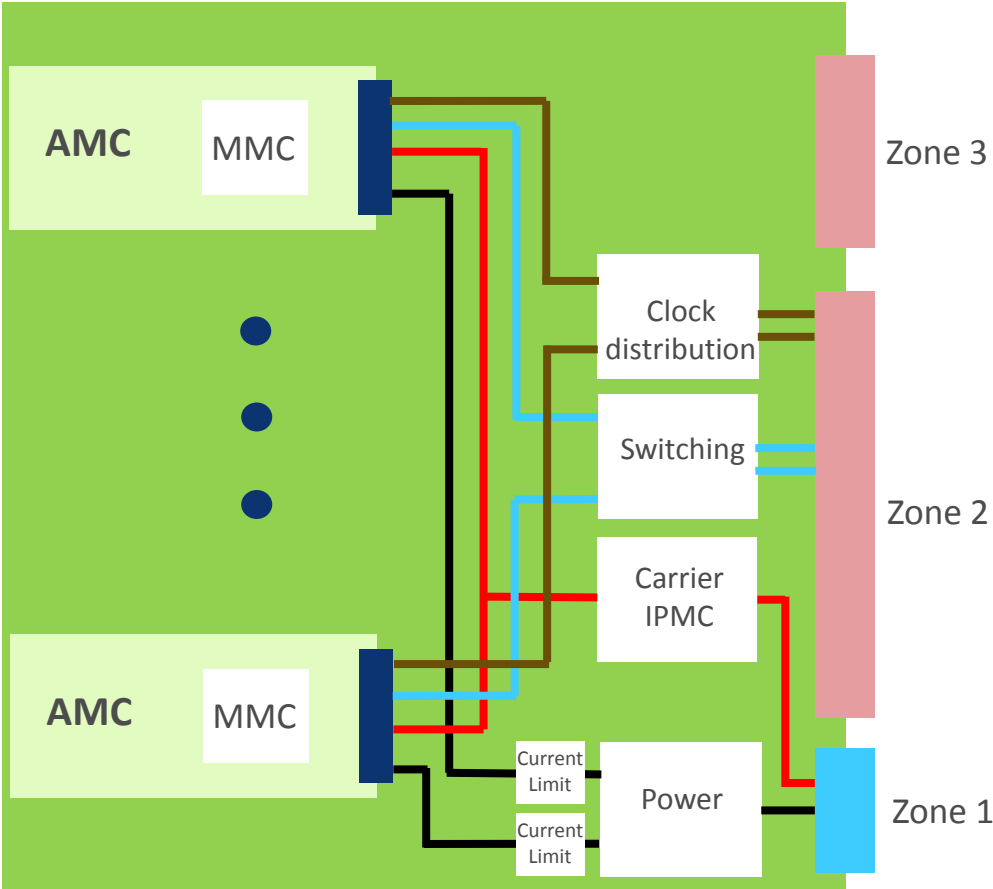


MTCA.0

- As MicroTCA does not use a Carrier board, the power, management, clock distribution and switching functionality must be realized onto another device
- Management Module: MCH (**MTCA Carrier Hub**)
 - IPMI management
 - clock distribution / generation
 - Switching functionality
 - JTAG slave / master
 - Redundant MCHs
- Power Module
 - 12V Payload Power
 - 3.3V Management Power
 - Redundant power modules
- Dedicated Slots for these modules are located in the MTCA Shelf

MTCA.0


ATCA Carrier



MTCA.0

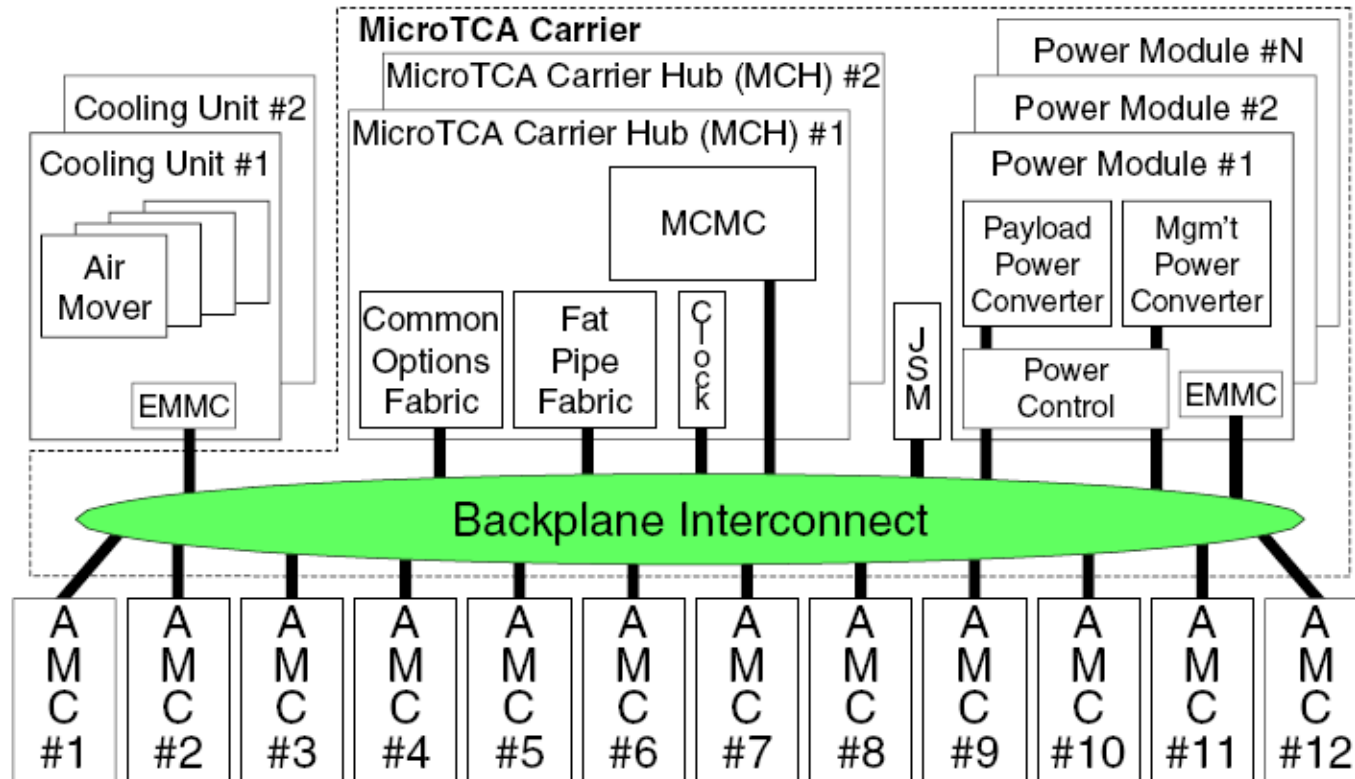
ATCA Carrier

MTCA Shelf



MTCA.0

MicroTCA block diagram



MTCA.0

Terms and Acronyms

- **MCH** MicroTCA Carrier Hub
 - This is the complete module you can buy from a vendor
 - **MCMC** MicroTCA Carrier Management Controller
 - This is the physical IPMI controller on the MCH
 - **MMC** Module Management Controller
 - This is the physical IPMI controller on an AMC
 - **EMMC** Enhanced MicroTCA Carrier Management Controller
 - This is the physical IPMI controller on a Cooling Unit and on Power Module
 - **IPMB-0** Intelligent Platform Management Bus 0
 - Logical IPMB, physically divided into redundant IPMB-A and IPMB-B
 - **IPMB-L** IPMB-Local
 - IPMI link between MCH and AMCs
-

MTCA.4

MTCA.4

Why were extensions needed to the existing MicroTCA specifications?

- **No Rear Transition Module (RTM) defined for MicroTCA**
 - Physics applications typically require a large number of I/O cables. It makes sense to connect them to the rear of the chassis

- **Special clock and trigger topology**
 - MicroTCA.0 specifies 3 Clocks and AMC.0 R2.0 specifies 4 Telecom and 1 Fabric Clock on the AMC Module. Physics applications typically need additional Clocks and Triggers

- **Sophisticated requirements for the clock and trigger accuracy**
 - MicroTCA / AMC defines typical telecom clock signals corresponding to PCIe values. Trigger signals are not specified

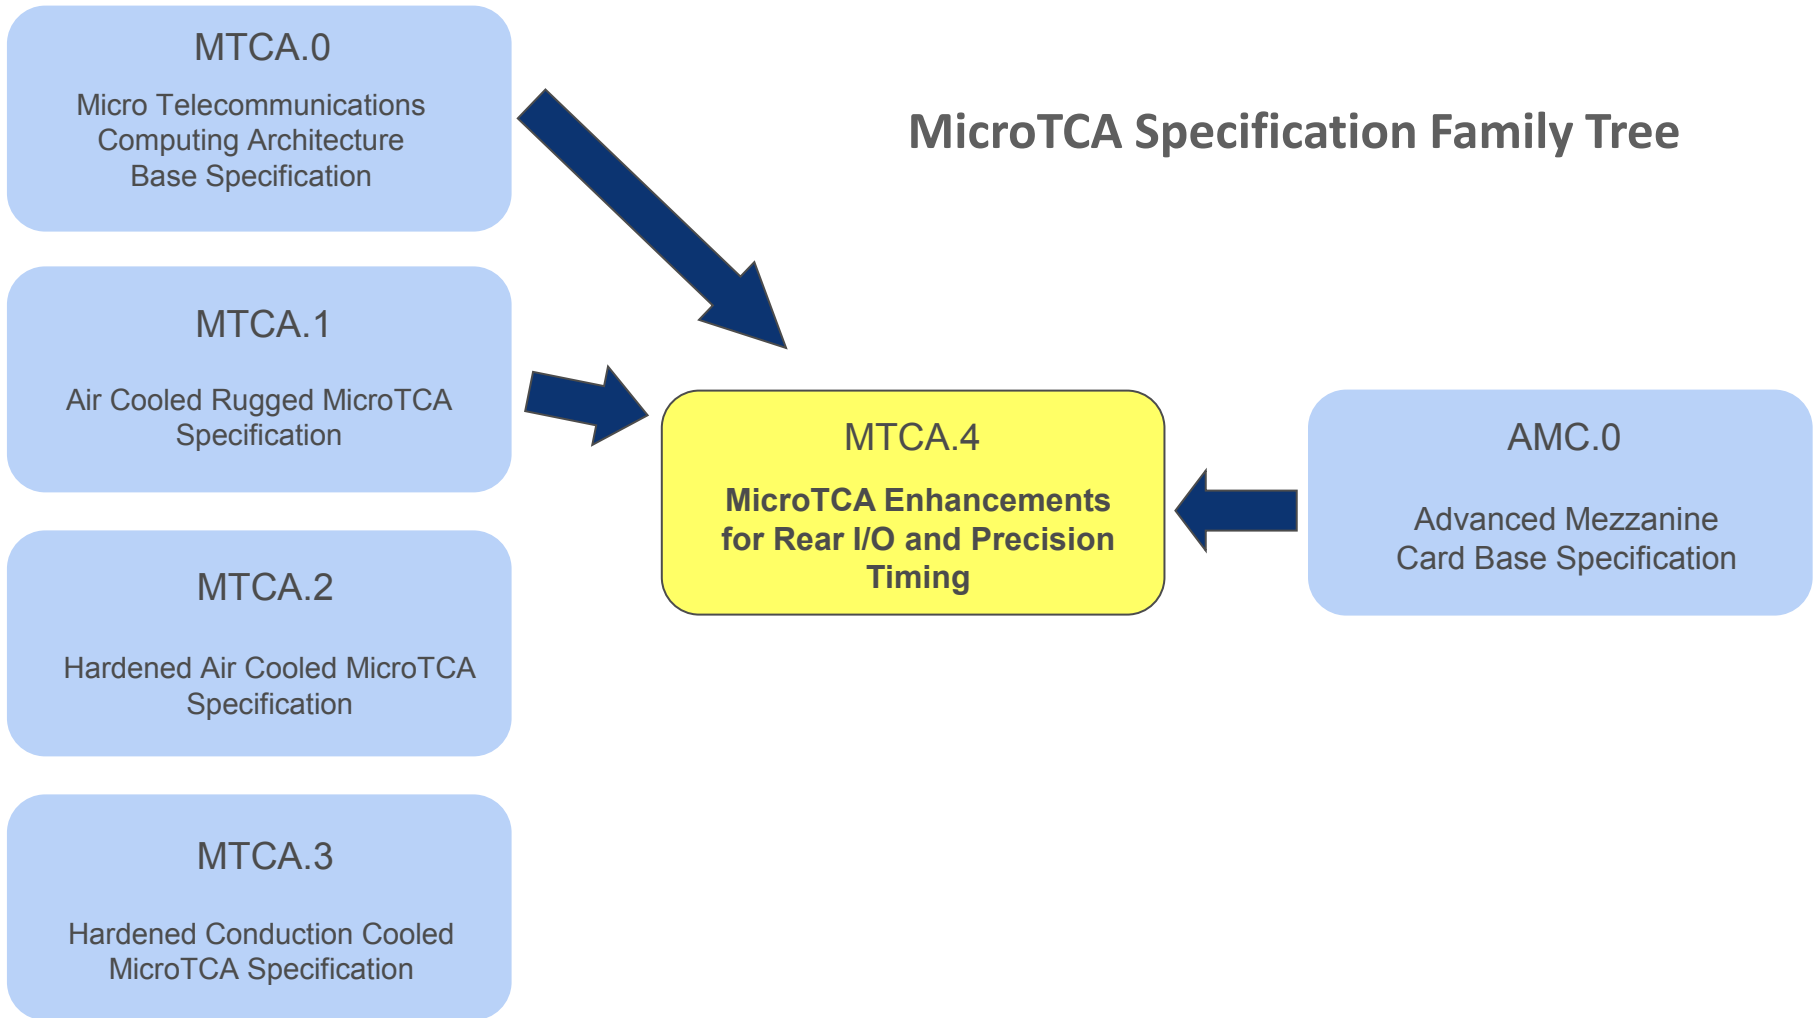
MTCA.4

Requirements for mechanics and sizes

- **AMC Module size: Double, Mid-size**
 - Allows for the max number of 12 AMCs in a 19" wide shelf
 - **Large MicroRTM real estate**
 - MicroRTM size approximately the size of the AMC (doubles depth of existing uTCA chassis)
 - **Use front panel mechanics based on Rugged MicroTCA (MTCA.1)**
 - Need to mechanically attach a module to avoid it being pushed-out by the corresponding module
 - Use Rugged MicroTCA retention device
 - **Reuse existing AMC front panels for the MicroRTM**
 - **Allowing mounting of mezzanine modules on the rear of the backplane**
 - **Optional zone 3 backplane**
 - **Define the management of the system**
 - **Suggest clocking and backplane topology**
-


MTCA.4

MicroTCA Specification Family Tree




MTCA.4

Features of MTCA 4 shelf with MicroRTM side view



MTCA.4

Module Sizes



MTCA.4

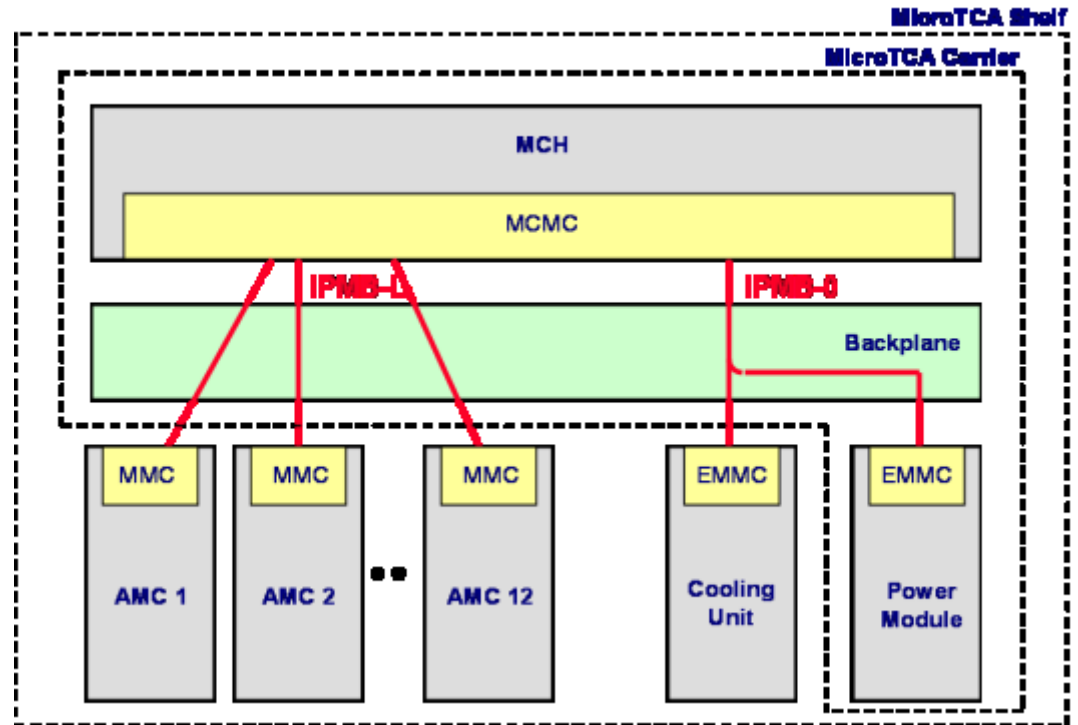
Management defined in AMC.0 / MTCA.0

- **IPMB-L**

- Connects the MCMC on the MCH to the MMC on the AMC Modules
- Radial architecture

- **IPMB-0**

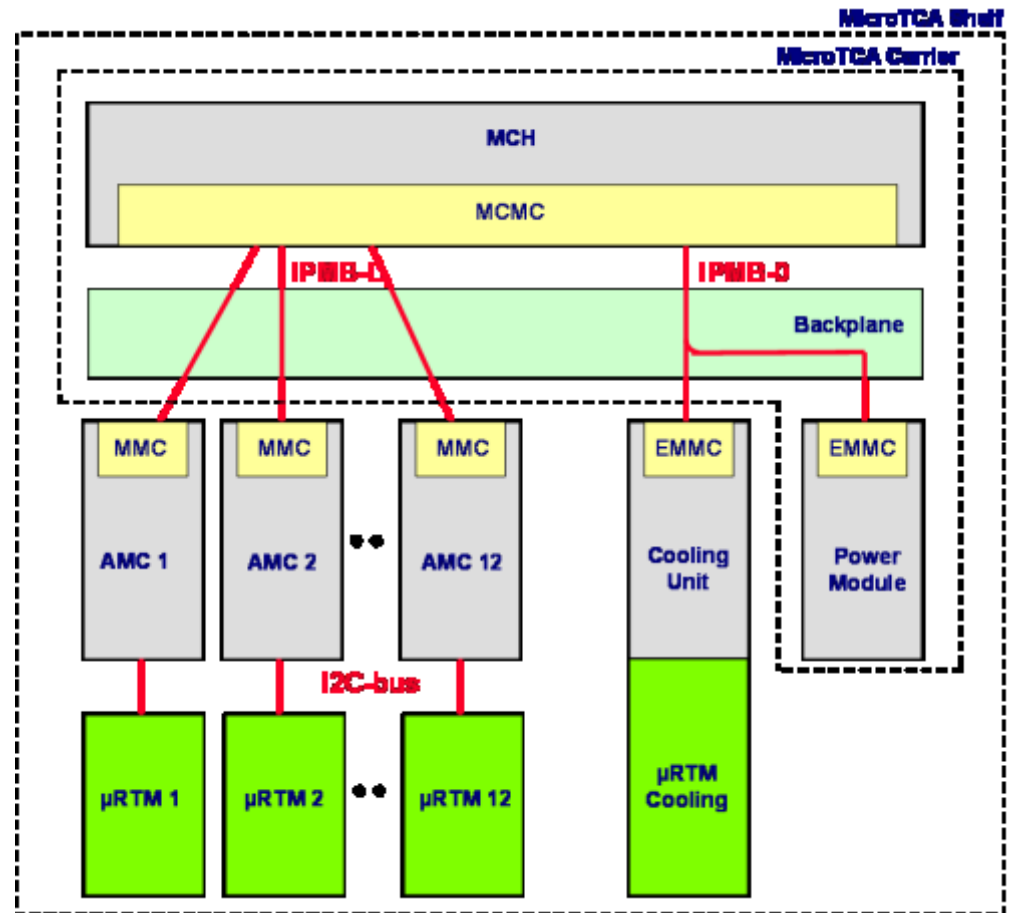
- Connects the MCMC on the MCH to the EMMC on the PM and CU
- Bused architecture



MTCA.4

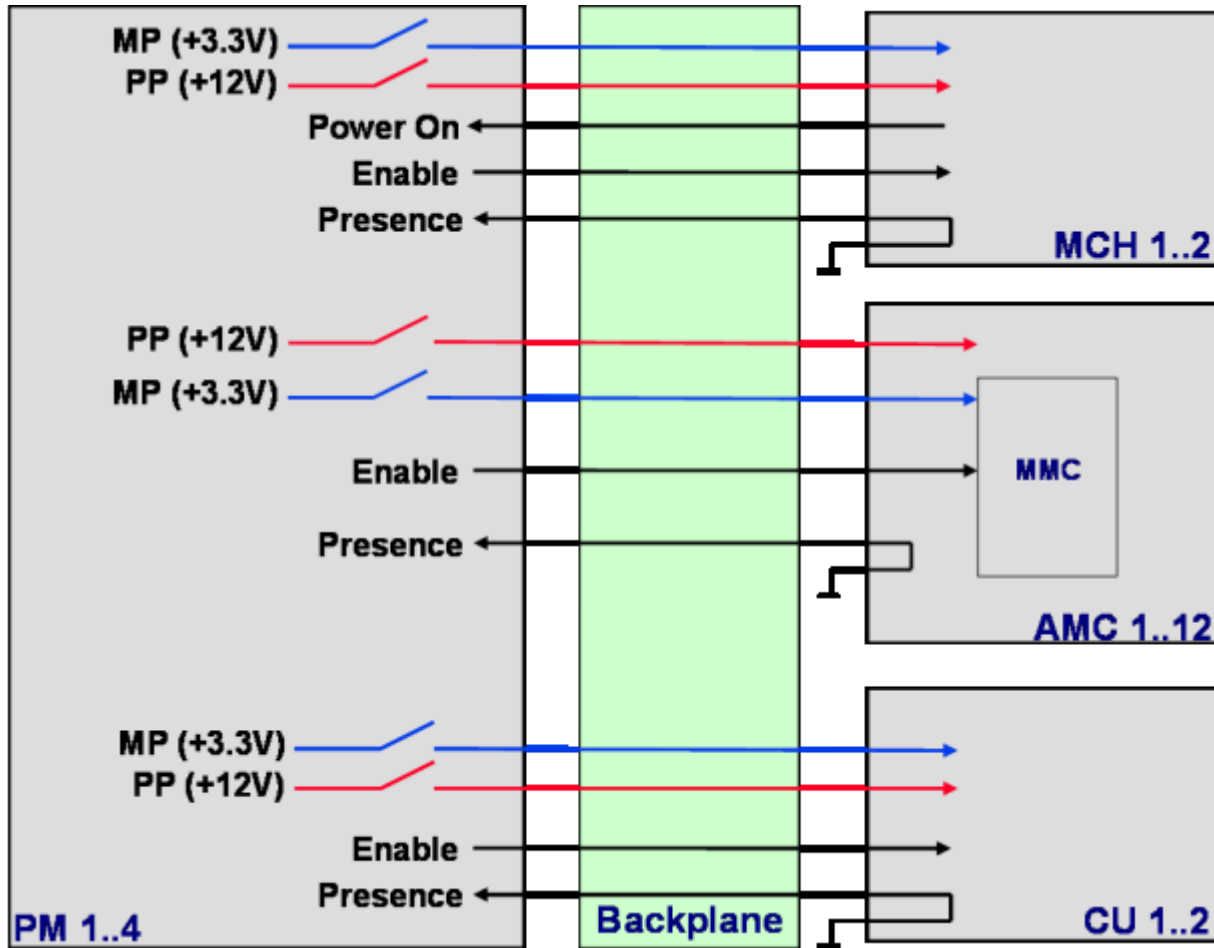
Management extensions in MTCA.4

- **IPMB-L**
 - Connects the MCMC on the MCH to the MMC on the AMC Modules
 - Radial architecture
- **IPMB-0**
 - Connects the MCMC on the MCH to the EMMC on the PM and CU
 - Bused architecture
- **I2C-Bus**
 - Connects the AMC to the μ RTM
 - The μ RTM is treated as managed FRU of the AMC



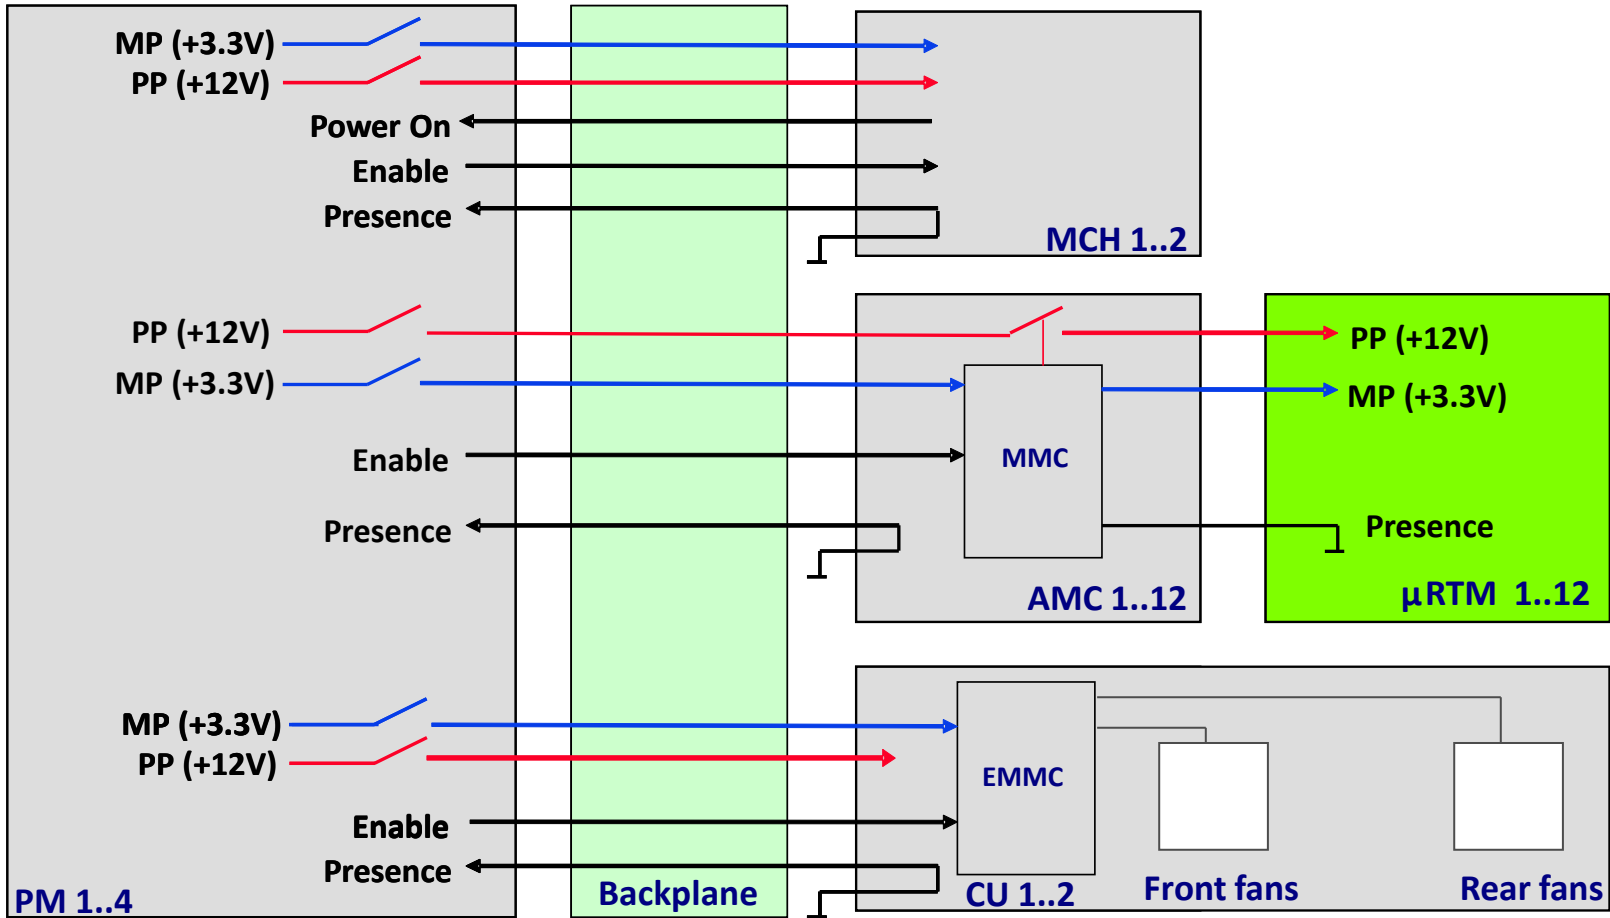
MTCA.4

Control signals as defined per AMC.0 / MTCA.0



MTCA.4



















Additional RTM control signals for MTCA.4




MTCA.4

Alignment and Keying

- mechanical keying prevents a module from being inserted which is not electrically compatible and could cause damage
- Eight keying positions are implemented that define the electrical interface

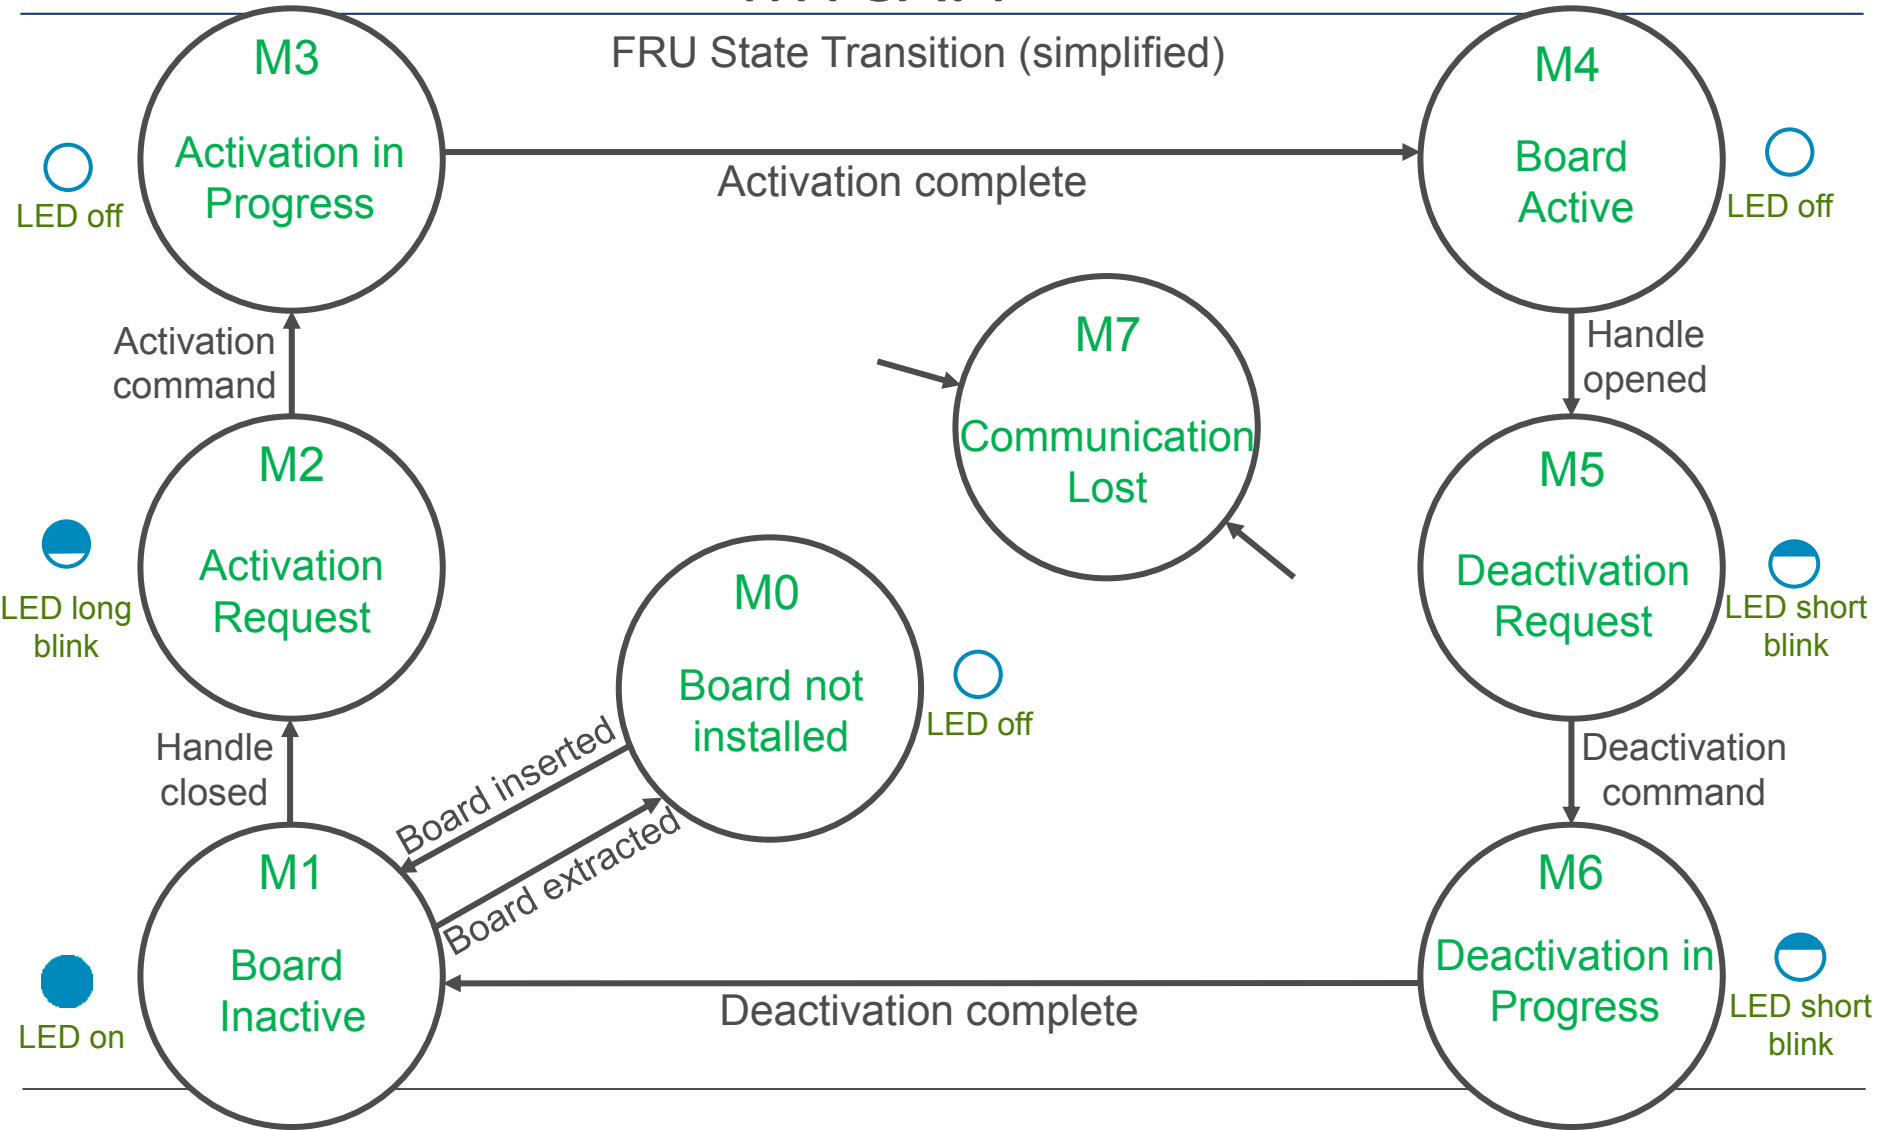
N	A Rotation in degrees	View into rear of AMC	View into rear of μ RTM
		Receptacle	Post
1	0		
2	45		
3	90		
4	135		
5	180		
6	225		
7	270		
8	315		
0	NA		

N	Data Signal in Volts
1	LVDS
2	0 - ± 1
3	$> \pm 1 - \pm 3.3$
4	$> \pm 3.3 - \pm 10$
5	$> \pm 10$
6	Reserved
7	Reserved
8	Reserved



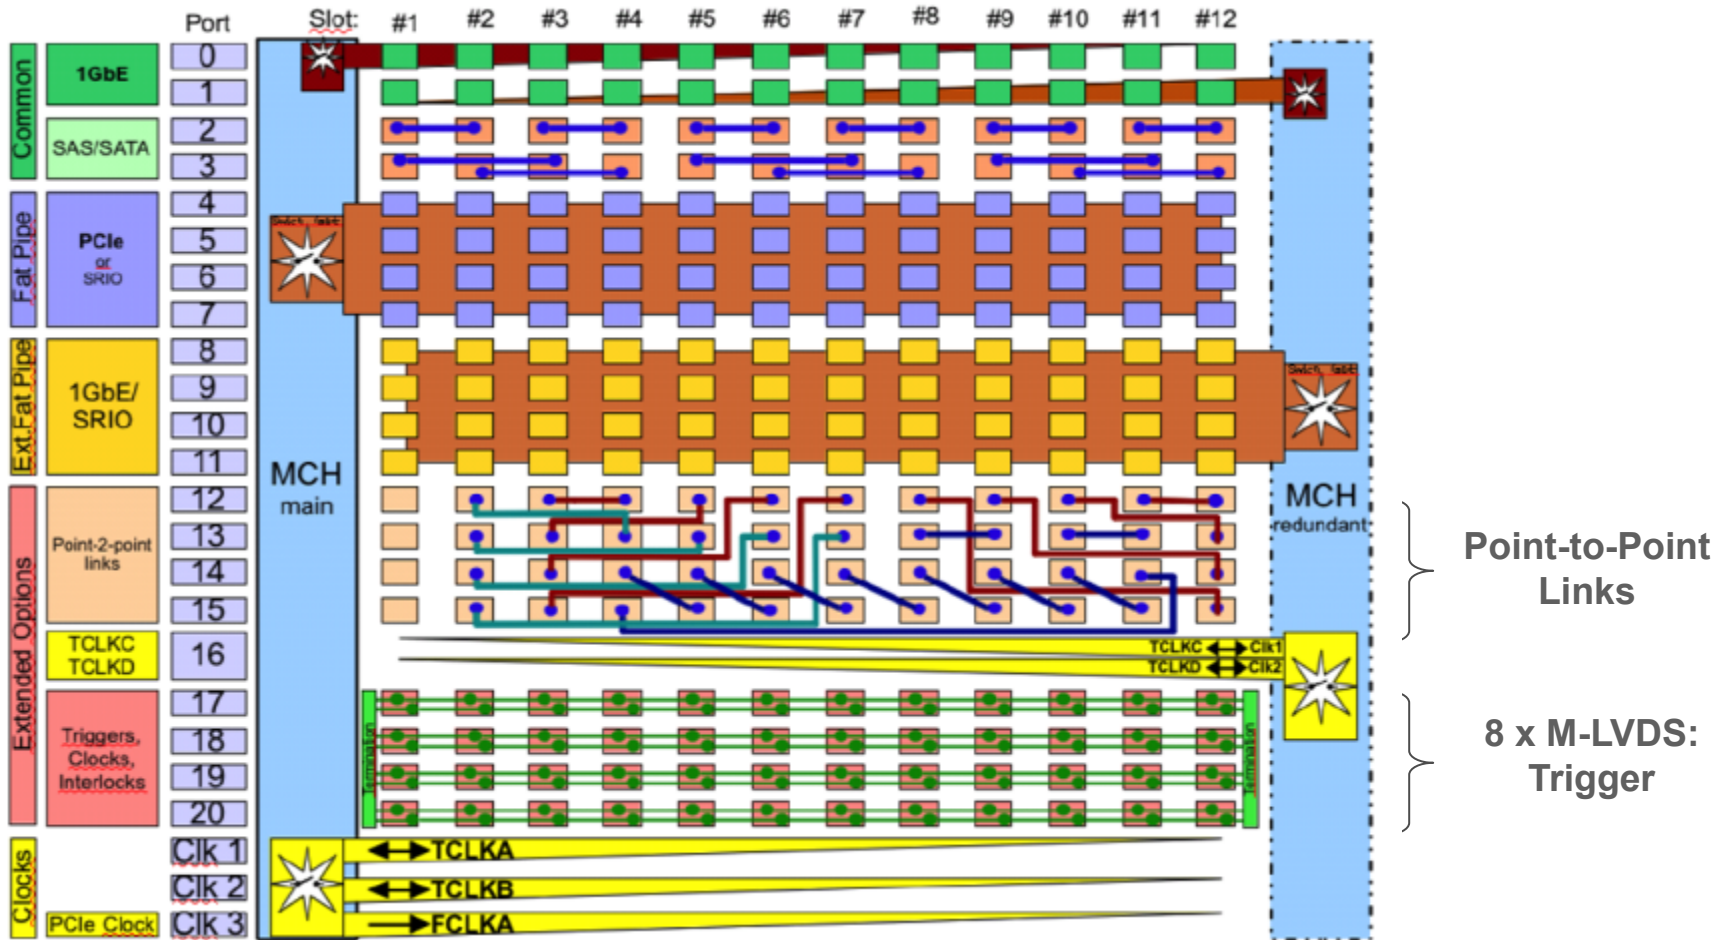
MTCA.4

FRU State Transition (simplified)



MTCA.4

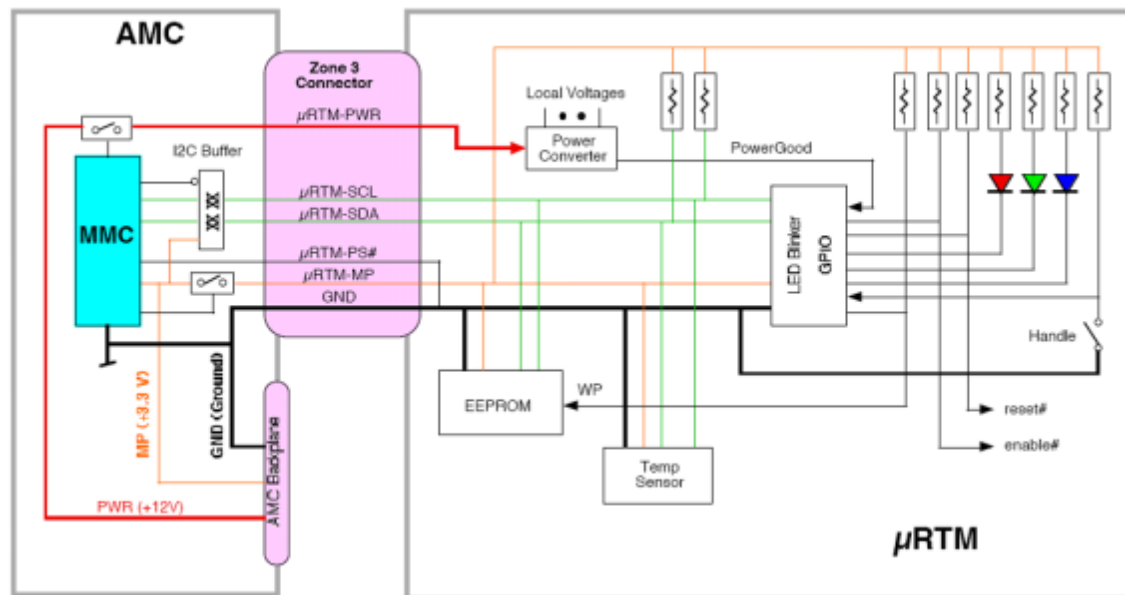
12-Slot MTCA.4 Backplane



MTCA.4

MicroRTM Management


- A management interface is defined on the lower zone 3 connector
- Management and power signals:
 - μ RTM-MP: Management Power for the EEPROM, Temp. Sensor and I/O Expander
 - μ RTM-PWR: Payload power for the RTM
 - μ RTM-PS#: RTM Presence signal, grounded on the RTM
 - μ RTM-SCL/SDR: I²C bus coming from the AMC MMC going to the RTM



MTCA.4

Front board and MicroRTM power distribution

- The total power for a slot (front board and RTM) is supplied through the front board AMC connector
- The MicroRTM power is supplied from the front board through the Zone 3 connectors
- Total available power for a slot is 80 Watts, the MicroRTM power is limited to 30 Watts
- The power required by the MicroRTM is subtracted from the power for the front board



MTCA.4

Shelf Cooling

- Some Physics applications require independent cooling of the front boards and the rear boards due to thermally sensitive oscillators on the rear boards. Also, independent fan speed control contributes to lower noise and power consumption
- The fans for the front board and the rear board are on one physical Cooling Unit
- Existing specifications did not deal with independent fan speeds on one physical Cooling Unit
- MTCA.4 enhances the commands „Get Fan Speed Properties“ and „Get/Set Fan Level“ to include a bit telling the MCH whether the Cooling Unit supports independent cooling and telling the Cooling Unit whether a „Get/Set Fan Level“-command targets the front or the rear fans

MTCA.4

MTCA.4 fan speed control commands

Table 3-9: MTCA.4 Get Fan Speed Properties

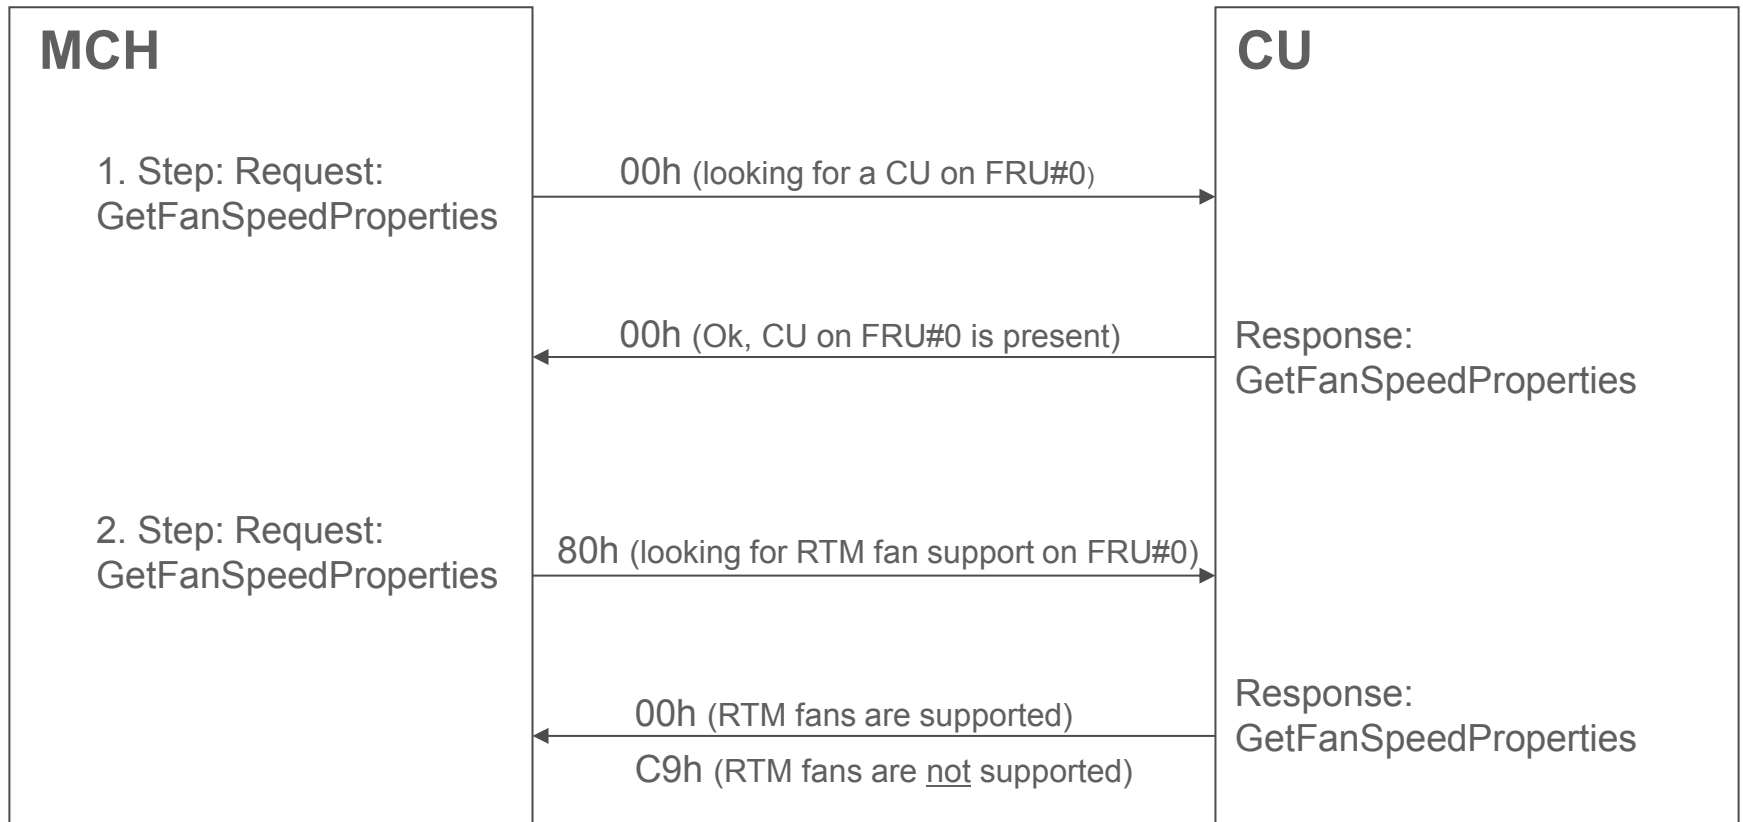
Request sent by the MCH

Response sent by the CU

Data Type	Byte	Data Field
Request Data	1	<i>PICMG Identifier</i> . Indicates that this is a PICMG-defined group extension command. A value of 00h is used.
	2	If the Fan Tray does <i>not</i> support RTM Fan control: <i>FRU Device ID</i> . Indicates the FRU device for which the command is intended. If the Fan Tray <i>supports</i> RTM Fan control: [7] – <i>RTM Fan Speed Properties</i> . This bit is set to 1b to indicate that the Get Fan Speed Properties query is for the RTM fans. This bit is set to 0b to indicate that the Get Fan Speed Properties query is for the front fans. [6..0] – <i>FRU Device ID</i> . Indicates the FRU device for which the command is intended.
Response Data	1	<i>Completion Code</i>
	2	<i>PICMG Identifier</i> . Indicates that this is a PICMG-defined group extension command. A value of 00h is used.
	3	<i>Minimum Speed Level</i> . This field describes the minimum setting that is accepted by the <i>Set Fan Level</i> command.
	4	<i>Maximum Speed Level</i> . This field describes the maximum setting that is accepted by the <i>Set Fan Level</i> command.
	5	<i>Normal Operating Level</i> . This field represents the default normal fan speed recommended by the fan manufacturer.
	6	<i>Fan Tray Properties</i> . This field holds properties of the Fan Tray. [7] – <i>Local Control Mode Supported</i> . This bit is set to 1b if the Fan Tray supports automatic adjustment of the fan speed. [6:0] – <i>Reserved</i>

MTCA.4

Sequence of IPMI commands to determine whether the CU supports independent MicroRTM cooling



MTCA.4

Air flow measurements in MTCA.4 Shelves

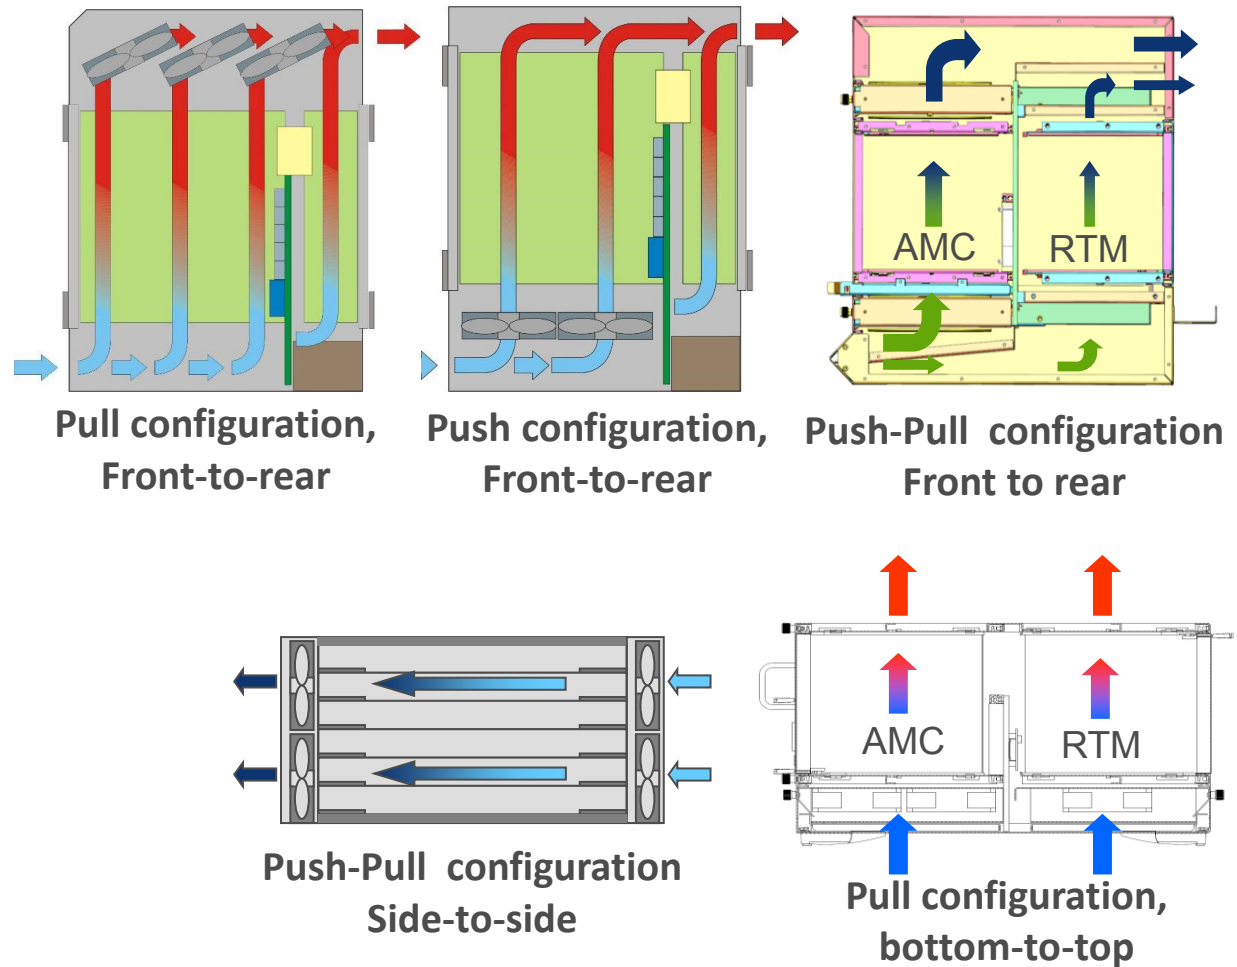
- One of the most critical issues in a shelf is cooling of the installed modules
- A reliable method to define the cooling capability of a shelf is to measure the volumetric air flow in m^3/h or cfm
- Measured Air flow values:
 - Bulk air flow value for the total shelf air flow
 - Individual air flow per slot
- Cooling capacity can be calculated based on the Air Flow and desired temperature differential between air intake and air exhaust

MTCA.4

Cooling concepts

The cooling concept depends on the installation situation of the chassis:

- Front-to-rear air flow
- Side-to-side air flow
- Bottom-to-top air flow
- Front-to-side air flow



Fan configuration:


- Push
- Pull
- Push-pull

MTCA.4

- Bulk air flow measurements in wind tunnel:




- Per slot air flow measured with air flow measurement boards and Flow Impedance Boards (similar to cp-ta in ATCA)




MTCA.4

- Test results:
Bulk air flow measured in wind tunnel = 635 m³/h
Sum of per slot air flow (front and rear) = 644 m³/h

Air distribution, front Slots, normal operation



Air distribution, Rear board, normal operation



- Cooling capability approximation: $Power = Air\ Flow * \Delta T / 3.3$
AMC slot 1: 24.3 m³/h => cooling capability (at $\Delta T=12k$) ~ 88 Watts
RTM slot 7: 12.1 m³/h => cooling capability (at $\Delta T=12k$) ~ 44 Watts

MTCA.4

How can MTCA.4 improve the reliability of the system?

- Prediction of failures (e.g. a fan does normally not fail instantly)
- Monitoring of every temperature sensor on every module and FRU
- Isolation of faulty modules(e.g. powering down over-heated modules)
- Power management
- Redundancy

MTCA.4

Redundancy

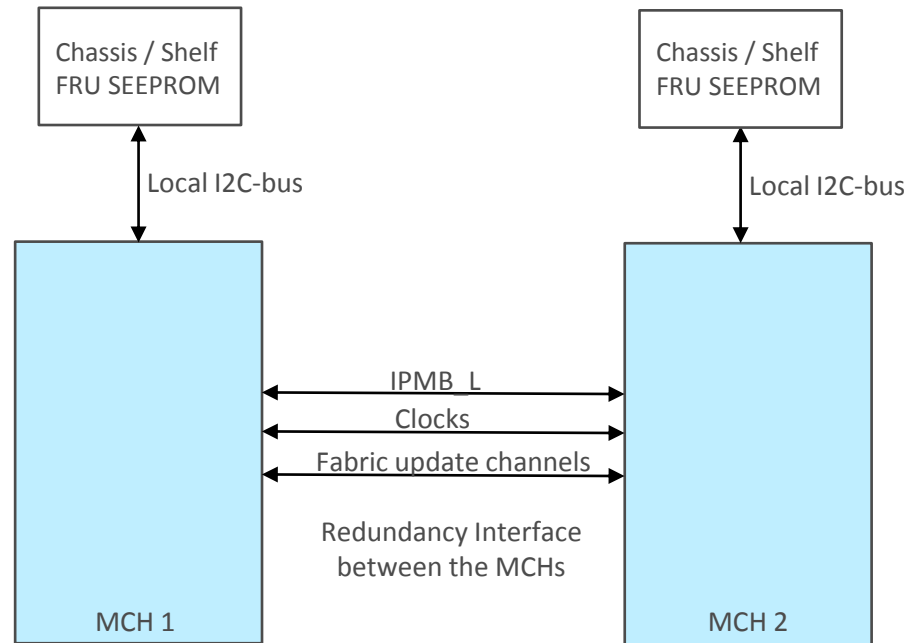
- For high availability applications all modules are redundant:
- 2 x MCH
- 4 x Power Module
- 2 x Cooling Unit
- IPMB-0:
One logical bus divided into two physical busses: IPMB-A and IPMB-B



MTCA.4

MCH Redundancy

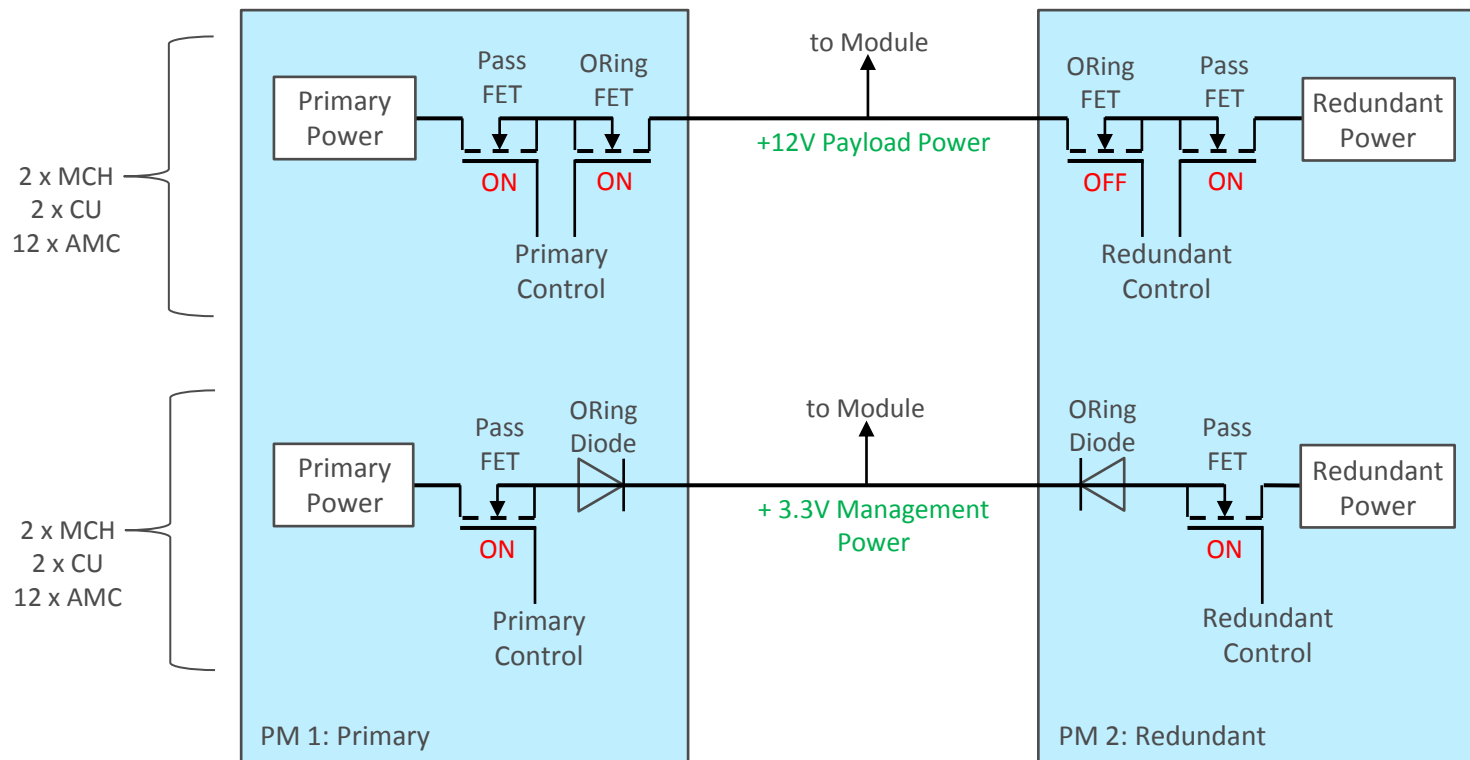
- Two MCH: One is Master, One is Redundant
- Redundant chassis / shelf FRU Information SEEPROM
- Redundancy Interface between the two MCH
- Redundancy defined in chassis / shelf FRU information



MTCA.4

Power Module Redundancy

- Up to 4 Power Modules per chassis
- Redundancy mode defined in shelf FRU file
- Individual power channel to each module and FRU



MTCA.4

Cooling Unit Redundancy


- Redundant Cooling Units in push-pull configuration
- Scenario 1: fan failure
- Scenario 2: Cooling Unit replacement

- 1 – 6 Fan
- 7 Hot Swap bush button
- 8 CU1
- 9 CU2

Pull



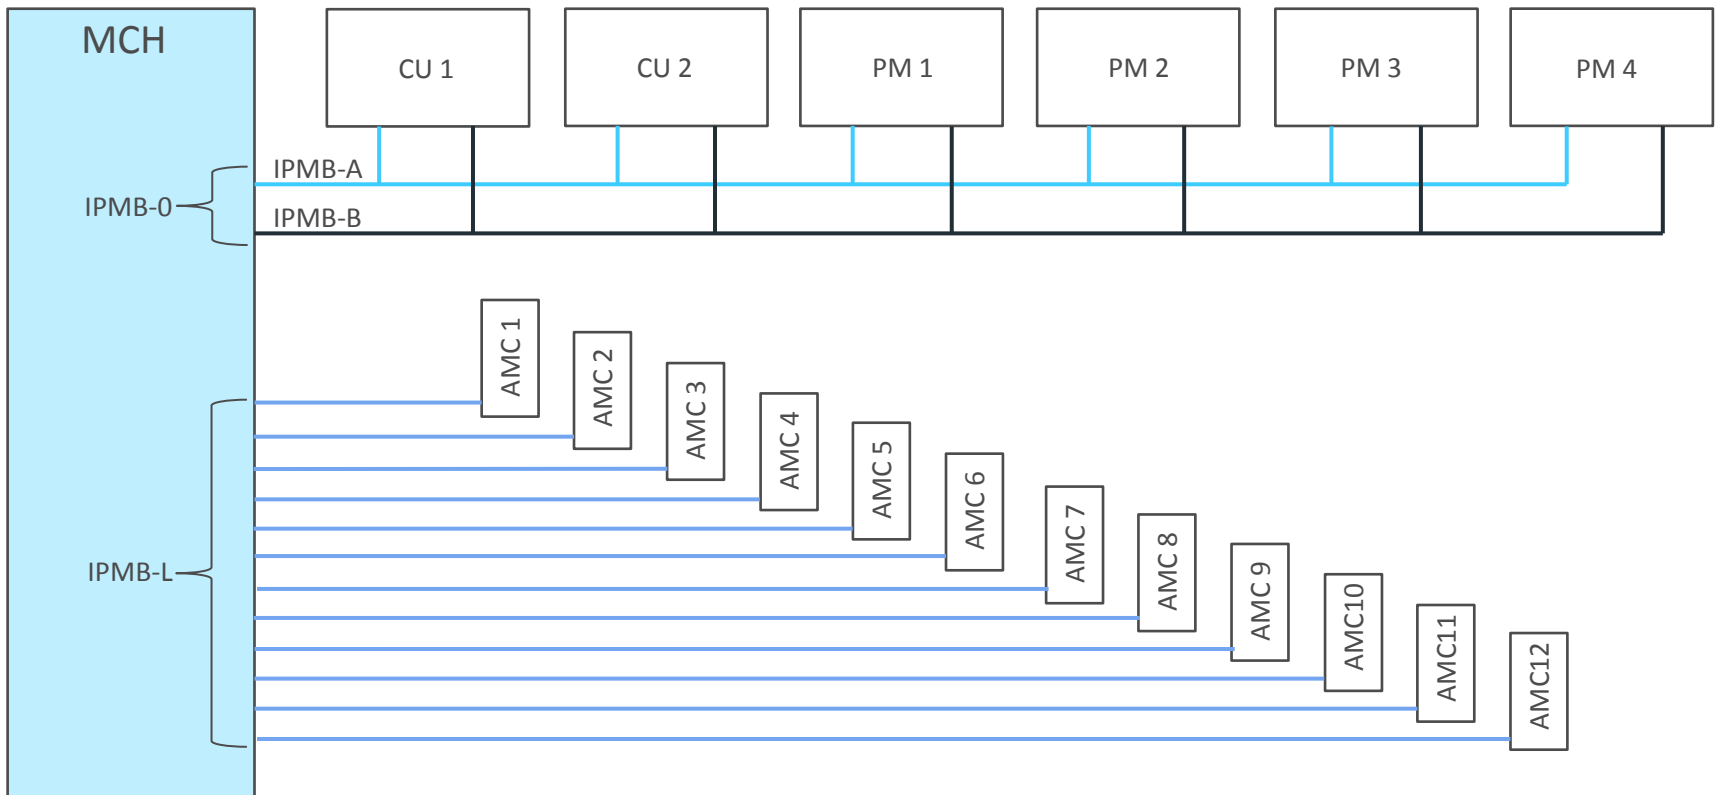
Push



MTCA.4

IPMB redundancy

- Individual IPMB-L to each AMC
- Redundant logical IPMB-0 to PMs and CUs



MTCA.4

MTCA.4 Chassis types

Various different MTCA.4 chassis available now:

- Laboratory use
- Fully redundant
- Compact sizes
- Small form factors




MTCA.4

Standardization continues: MTCA.4 Standardization activities


Definition of a RTM Auxiliary Backplane (MTCA.4.1)

- Based on the LLRF backplane
- Optional connector usage



Protective mechanical cover for AMC and RTM modules

- Protective cover to mechanically protect components
- For Side A and Side B



Thank you

Ralf Waldt

Pentair Technical Solutions GmbH

ralf.waldt@pentair.com