xTCA at CERN

an overview based on input from many kind colleagues

M. Joos, PH/ESE

ALICE

- Will use the AMC40 cards developed for LHCb as the basis of the new Common Read-out Units (CRU) for the read-out and concentration of data of those subdetectors that will upgrade during LS2. (Including TPC, TDR, MCH, ITS, etc.)
- The CRUs will be a common interface between the upgrading subdetectors and the new online system (DAQ)
- Most CRUs will be located in the counting rooms in ATCA crates. A PCI Express card CRU is also considered, however ALICE presently preferes CRUs detached from the DAQ computers (that is, AMC cards in ATCA crates)
- ALICE choose the GBT optical links as the front-end side links of the CRUs, while the server side links will be either 10 Gb/s Ethernet or PCIe over optical fiber.
    - Input and output links will be connected through the front panel
    - The carrier boards will have limited functionality
- At the present level of thinking, an AMC40 could have 24 front-end side GBT links, (while having 12 server side links), and 4 AMC cards can be placed on a ATCA carrier. If we calculate with a maximum density of 12 carriers per shelf, we will need the following number of shelves as a minumum:
    - TPC: ~ 6640 links 6 shelves
    - TRD: ~ 1044 links 1 shelf
    - MCH: ~ 500 links 1 c shelf (probably shared with MID, ZDC, etc.)
    - ITS (after concentration): 1 shelf
    - Trigger distribution 1 shelf in magnetic field (may be a custom crate)
- ~ 10 shelves in total
- Types and manufacturers have not been investigated yet
- Contact persons: alex.kluge@cern.ch (ALICE electonics co-ordinator), tivadar.kiss@cern.ch (CRU development project leader)

ATLAS - I

Project: Support for ATCA hubs (ATLAS – IT collaboration, Input from S. Batraneanu)

- Existing setups (shelves and hubs) are still immature and the requirements are vague
- Big differences between products coming from different manufacturers and the market is quite dynamic
    - Products of Kontron, Emerson and Vadatech have been looked at
- Cannot recommendation a hub product because none of them fits all known requirements

Project: New barrel Sector Logic boards (Riccardo Vari)
Target: Phase-2, Card design will not start before L2, Plan is to develop one ATCA blade

Project: New Small Wheel Trigger Processor (Lorne Levinson)
Target: Phase-1
Details: Use Liquid Argon boards or those from the SRS project. In all we expect 16 boards, each with mezzanines. Input is about 1700 fibres

Project: Double AMC board (+ commercial ATCA carrier) to readout the TileCal (Alberto Biot)
Target: Demonstrator at the end of the year
Status: Finishing the layout and will produce the first prototype in March

Project: L1Topo (David Sankey)
Target: LS1
Status: Prototype under test at CERN (B104), PRR Mar/April

ATLAS - II

Project: eFEX and jFEX systems (David Sankey)
Target: LS2
Components: eFEX module, jFEX module, Hub module, ROD FMC, FEX Tester Module (FTM)
Status: Prototype PDRs under way (PDR for FTM)
Concerns:

- Power / Cooling
  - Power dissipation in vertical airflow mode in standard ATLAS rack
- PCB issues (VIAs etc.) for (long) high speed serial links
- Integration into DCS

Wish list from external groups:

- Help with the evaluation of shelves with a cooling capacity of 400 W per slot. These shelves should be compatible with the environment in USA15
- Integration into DCS

Project: Fast tracker (Jinlong Zhang)
Target: LS2 (installation starts in 2015)
Components: Data formatter and FLIC (~34 blades in total)
Status: Prototype boards with minimal IPMC functionality but LAPP IPMC to be used for the final board
Concerns: Availability of IPMC

ATLAS - III

Project: LAr LDPS (John Hobbs, Stonybrook for Carrier blade, Guy Perrot, LAPP for AMC)
Target: LS2
Components: LDPB (Carrier Blade with AMCs), commercial 10GbE switch blade
Status: Prototypes being developed, expected end of year, Production for LS2
Concerns:

- Front panel optical connectors and uPods on AMC
- Fibers arrival and spare length storage in racks
- Shelf cooling issues (vertical versus horizontal cooling)
- Mix of PCs (FELIX) and shelves in a rack
- Blade and AMC cooling issues
- Integration into DCS

Wish list from external groups:

- Help with the evaluation of cooling in standard USA15 racks, should we want to use them
- Integration into DCS

Project: LAr Upgrade Demonstrator (Guy Perrot LAPP, Yuji Enari Tokyo)
Target: LS1
Components: ABBA (blade), commercial 10GbE switch blade
Status: Board available, firmware in development. One blade needed for readout in April 2014 at EMF. 2 blades for Summer in USA15
A few blades produced by LAPP & Tokyo
Concerns: Firmware development which takes time

ATLAS IPMC

- Mechanical: DDR3 form factor Mini-DIMM
- Based on 32-bit ARM Cortex M4 µC
- Features
  - IPMBus with on board buffers, H/W address detection
  - Hot Swap management with ATCA LEDs and front panel switch
  - Management of up to 8 AMC + RTM
  - On board Event LOG
  - FRU & SDR access via I2C
  - Access to ATCA board sensors via I2C
  - IPM_IO: Configurable User Signals for Payload management, e-keying
- Tested with a mix of Boundary Scan tests (internal connections) and operational tests (connectors)
- Coming soon: full Boundary Scan test
- Documentation: http://lappwiki.in2p3.fr/twiki/bin/view/AtlasLapp/ATCA
- 20 boards manufactured
- New production (V2.2) foreseen mid 2014 then yearly
  - Users have been contacted to get their quantities and date requirements
  - In total 220 pieces requested by ATLAS, LHCb and BELLE projects

Software: "ICARE"
- Written in standard ANSI C
- GCC (4.7.0) tool chain
- Configuration Management: CMT
- FRU (ATCA board) Hex generation utility (using M4 preprocessor)

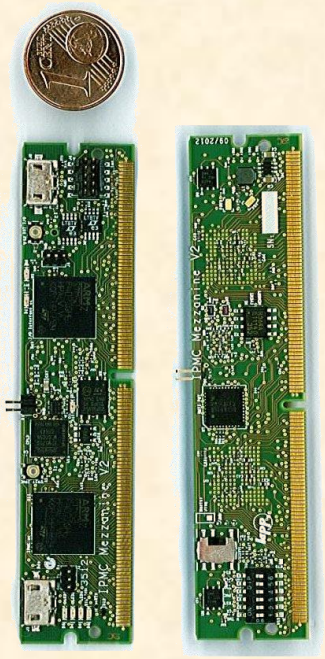

Contact: Guy Perrot
See also: https://indico.cern.ch/event/300897/

ATLAS – DCS Integration

Goal: Monitoring/control within WinCCOA (PVSS)

- Shelf monitoring: voltages, currents, temperature
- Shelf control: power (on/off/cycling)
- Board monitoring: be able to monitor selected board parameters
- Board control: reset individual boards, etc.

Data path:

- Access via SNMP
  - Pigeon Point shelf manager has built-in SNMP agent
  - All selected data should be made available in shelf manager (shelf and board data)
- Read/subscribe with existing WinCC OA SNMP client

➲ No need for custom interface software

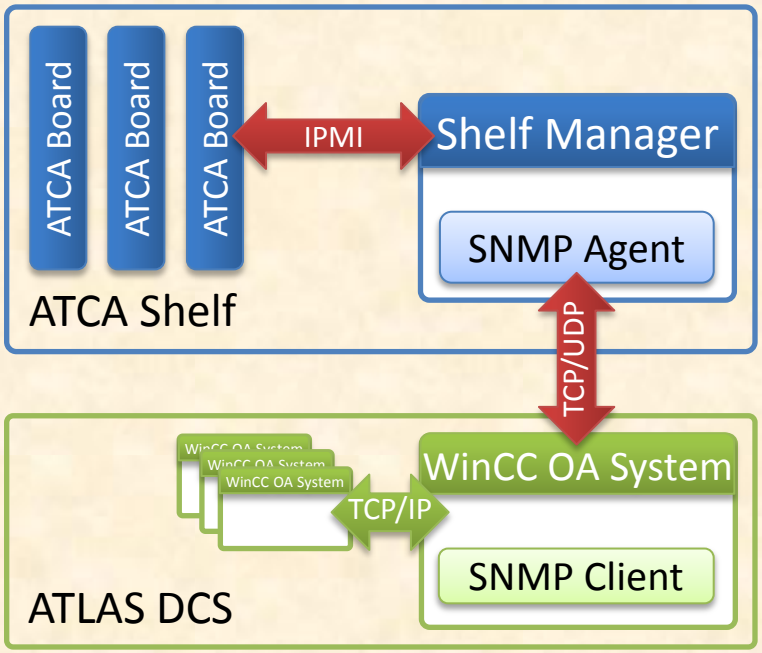

Status:

- Prototype set up (for ATLAS CSC readout upgrade): can read parameters and control power from WinCC OA (or other SNMP client)
- Performance and shelf manager load constraints evaluated
- Experiment network integration being discussed with IT, some protection required
- Missing experience with integration of parameters for actual ATCA boards
- To do: High-level WinCC OA integration into experiment framework

Contact: Stefan Schlenker

CMS

- Standardize key items in order to ease commissioning and long term maintenance of upgraded systems
  - Shelf
  - CMS interface, a.k.a. AMC13
    - Data Acquisition (DAQ)
    - Trigger Timing and Control (TTC)
    - Trigger Throttle (TTS)
  - AMC Module Management Controller (MMC)
    - UW design; very complete, used in e.g. CTP*, MTF*, uHTR, AMC13
    - UK design; also very complete, used in e.g. MP7, FC7
    - DESY/CPPM/CERN design; supported by CERN; includes mezzanine ref design
  - Register access software and firmware
    - IPBUS; includes a large software package and firmware
  - CMS DCS integration

Information provided by M. Hansen

CMS: Proposed uTCA Shelf

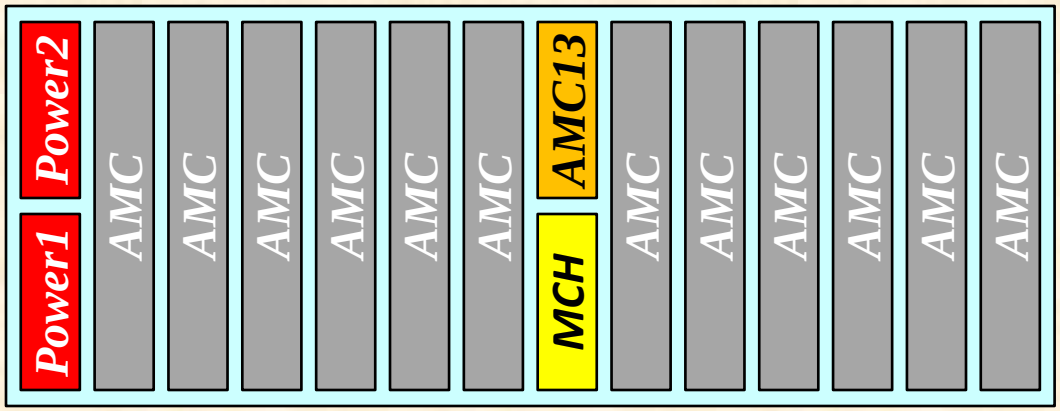

e.g. Vadatech VT892
~20 in CMS hands for evaluation

- 12 full size double height slots
- Vertical air flow
- Redundant power supply (optional)
- No RTM

- Redundant telecom backplane
  - Fabric C and D routed to MCH 1 & 2 tongue 2
  - Fabric E and F routed to MCH 1 & 2 tongue 3 & 4
  - Customized with additional backplane interconnects (VT894)

E. Hazen, BU

CMS: Planned uTCA Port Use

Note: Interconnections can be customized by the backplane manufacturer inexpensively.

Fabric	AMC Port	MCH1	MCH2	Category	MCH tounge	CMS Use
A	0	X		Common Options	1	GbE for slow control
	1		X			DAQ / Controls
B	2	X			2	Reserved
	3		X			Synch. Controls, TTS
Clock	TCLKA	CLK1/2		Clocks		Reserved
	FCLKA		CLK1/2			LHC Clock
D-G	4-7	X		Fat Pipes	3, 4	User
	8-11		X			
H-K	12-15			Extended Fat Pipes		
	16-19					

E. Hazen, BU

CMS: Interface

- AMC13
  - Tongue 1
    - TTC interface
    - AMC DAQ links over Fabric A
    - Three 5/10 Gb DAQ links
  - Tongue 2
    - Distributes LHC clock and ctrls
      - FCLKA and Fabric B
  - Tongue 3–Tongue 4
    - Service interfaces (base config)
    - Optional (TCDS)
      - TTCMi clock
      - TTCMi orbit
      - Two level 1 triggers (LVTTL)
      - Two auxiliary (NIM) inputs
    - Tongue 3 & Tongue 4 can be customized

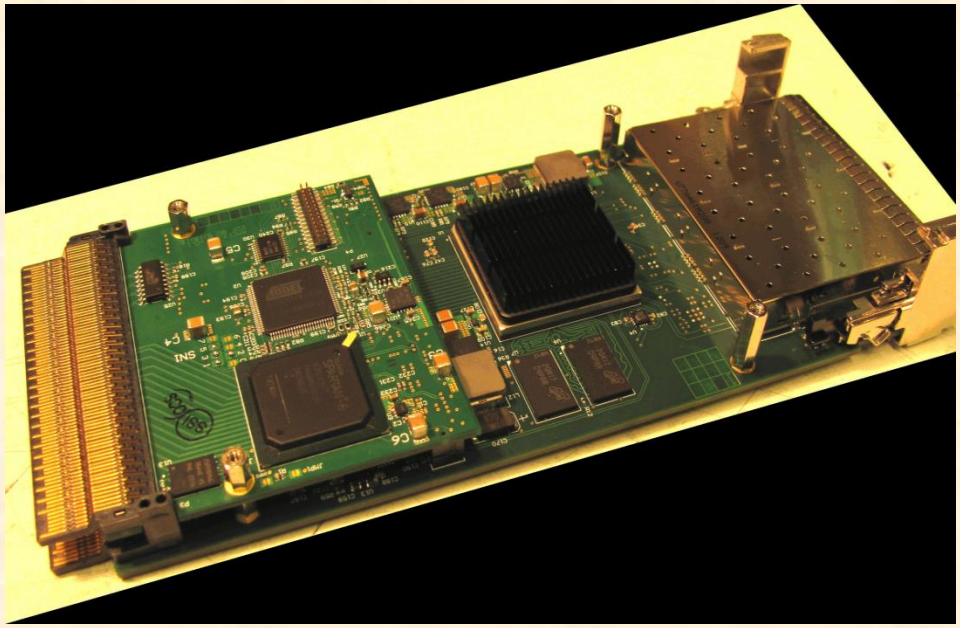

CMS Production ready modules

- uHTR (h-cal trigger and read out)
  - FE: XC6VLX240T
  - BE: XC6VLX195T
  - 24 Rx @ 6.4Gb (4.8Gb)
  - 12 Tx @ 6.4Gb (4.8Gb)
  - 2 TRx @ 4.8Gb
  - Ports 4,5,8,9 populated

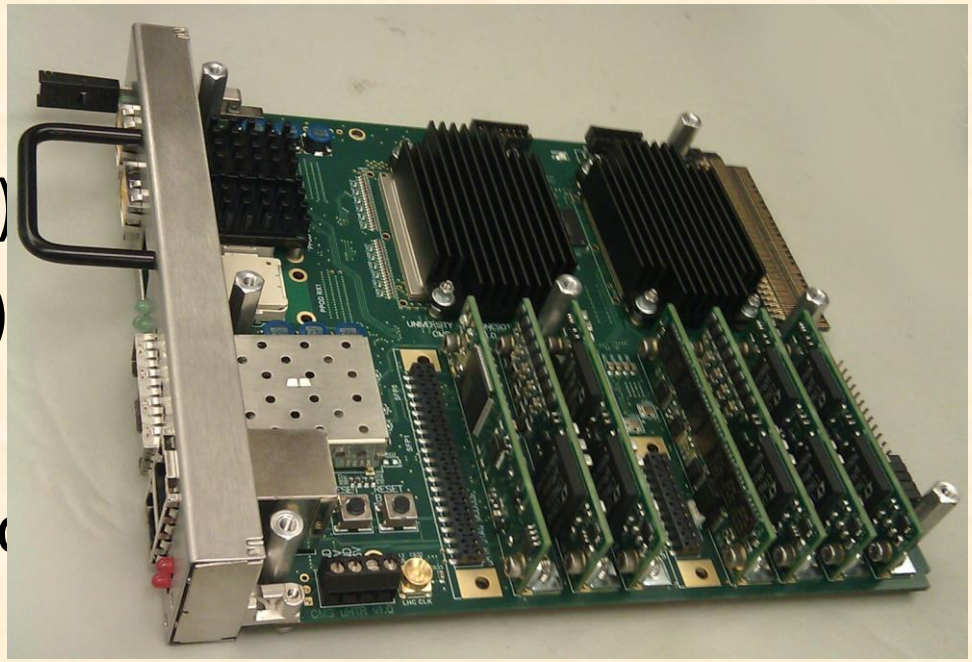

CMS Production ready modules

- MP7
  - XC7VX690T
  - 72 Rx @ 13Gb
  - 72 Tx @ 13Gb
  - Ports 4-8 populated
    - 9-11 LVDS

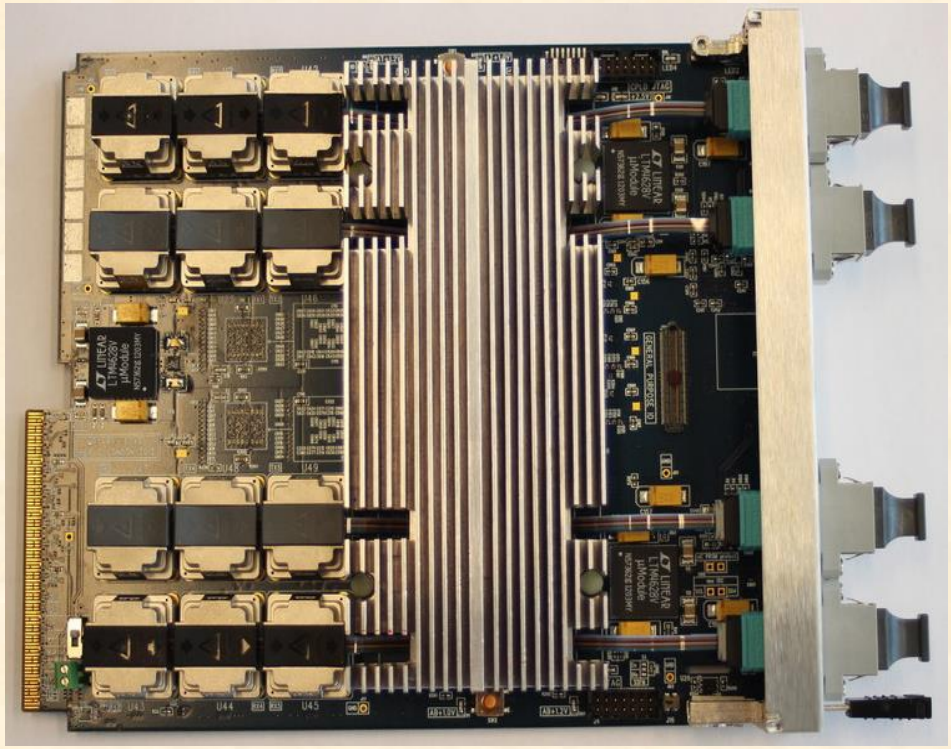

CMS Production ready modules

- FC7 (multi-purpose FMC carrier)
  - XC7K420T
  - Two LPC FMC sites
    - Site 1: 8 TRx @ 10Gb
    - Site 2: 12 TRx @ 10Gb
  - Ports 4-11 populated

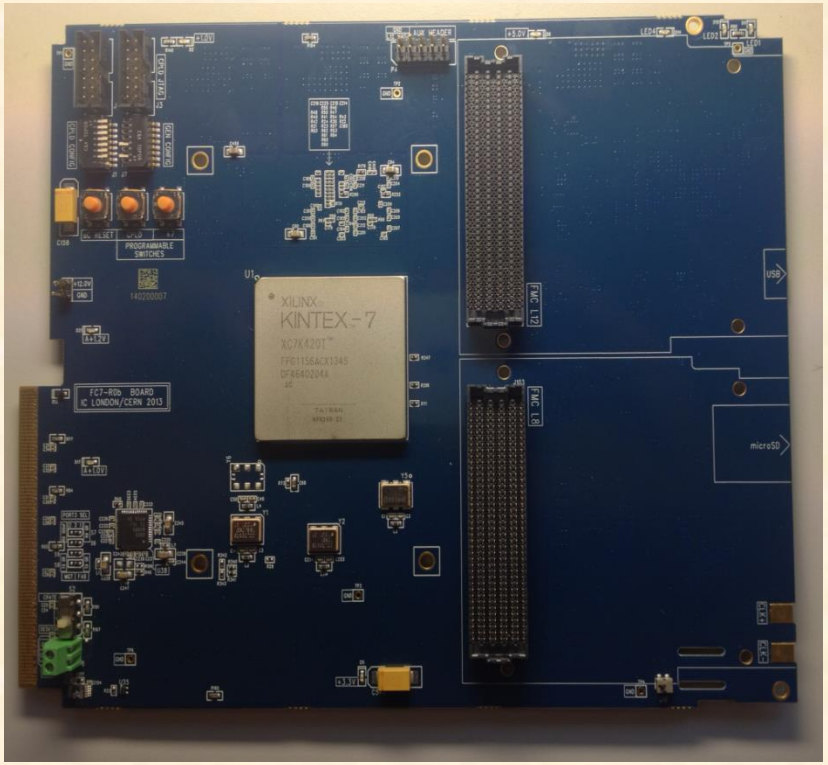

CMS Pre-production ready modules

- CTP7 (calorimeter trigger processor)
  - XC7VX690T
  - XC7Z045
  - 67 Rx @ 10Gb
  - 48 Tx @ 10Gb
  - Ports 4-7, 12-15, 17-20 populated (as VT894)

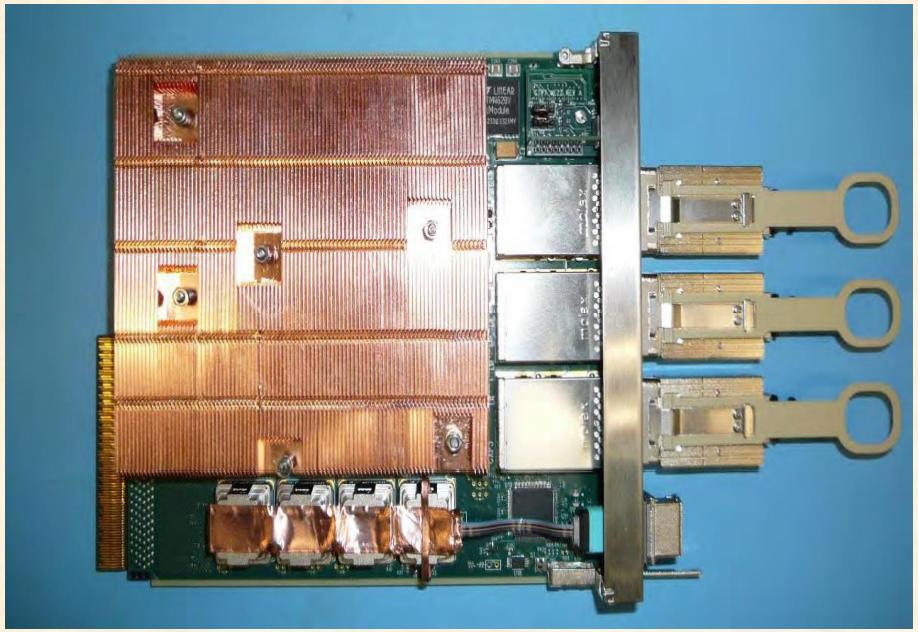

CMS Pre-production ready modules

- MTP7 (muon read-out)
  - Dual card stack
    - Occupies two slots
  - XC7VX690T
  - XC7K70T
  - 80+4 Rx @ 10Gb
  - 28 Tx @ 10Gb
  - Pt LUT module
    - Up to 2GB low latency RAM
    - Split in banks

CMS Commercial Items

uTCA crate
Vadatech*

Power-One AC/DC (48V) converter

Power Module
NAT or Vadatech

MCH
NAT or Vadatech

* Development with ELMA and SCHROFF ongoing to provide second source

LHCb

Contact: Ken Wyllie, Jean-Pierre Cachemiche

Global "xTCA strategy"

- ATCA pursued as the common format for off-detector electronics until now
  - successful hardware development and production (AMC40)
- R&D campaign launched recently to look at the possibility of moving the functionality to a PCIe format
  - PCI40 cards would be housed in PCs which form part of the DAQ & event-building network
  - Seems to have a major cost advantage
  - Review was held in February to assess the proposal of adopting PCI as the base-line with ATCA as back-up
    - Decision expected in March
    - If decision is to go for PCIe then ATCA will be completely eliminated (also in niche applications)

In what areas would xCTA be used

- Philosophy is to build generic hardware to cover all the needs of data-readout, timing & fast controls and slow-controls
  - Differentiation via F/W and I/O interfaces

What projects are currently under development

- ATCA40 carrier and AMC40

Deployment schedule for xTCA

- Prototypes debugged, kept as back-up solution

Main problems

- Power consumption (may be up to 350 W per blade)
- IPMI management (developing an IPMC costs a lot of time)

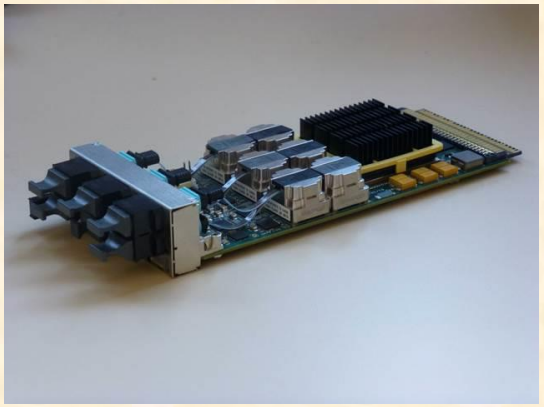

AMC40

What is on the "wish list"

- Common crate definition
- Common supervision software and hardware (IPMI)
- Tested inter-operable products (shelves, AC/DC)
- Definition of "LHC profiles". Examples:
  - Specific purpose data paths in ATCA mesh backplane interconnections (Clocks, data results, synchronizations, etc …)
  - Minimum set of signals to manage for interoperability (slow control, GBE, PCIe, clock distribution)

Activities of PH/ESE

xCTA evaluation project

Project definition (short version):

Evaluate COTS items of general interest and keep in contact with the (LHC) experiments

MTCA results and plans:

- Shelves
  - Cooling performance bad in some shelves
  - Issues with backplane alignment (can cause damage to AMCs) in some shelves
  - Different opinions in the experiments on:
    - Preferred AMC size (mid / full)
    - Need for RTM slots
- Power modules
  - Lots of trouble with modules from several manufacturers
  - No clear user preference (yet) for AC or DC input
- Interoperability
  - Lots of problems when combining components of different manufacturers
    - E.g. No stable configuration for segmented powering schemes
- Proposal for a "MCTA.CERN" system:
  - Vertical air flow
  - 12 slots for full size / double width AMCs
  - 6 RTM slots
  - 2 MCH slots
  - 2 (front) + 4 (rear) power module slots

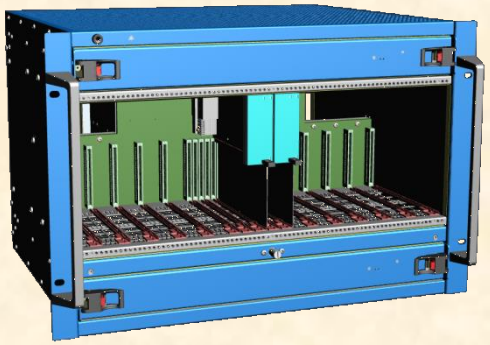

"MCTA.CERN"

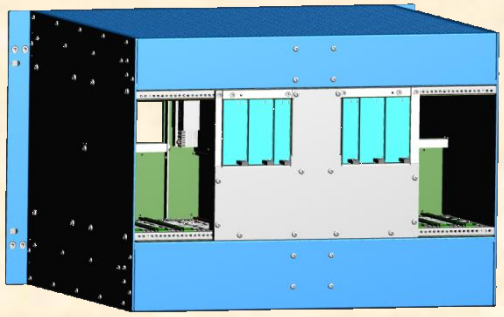

Activities of PH/ESE

ATCA results and plans:

- Evaluation focusses on:
  - Shelves
    - Cooling
      - Observation: Cooling performance bad in some shelves
      - So far only evaluated shelves with horizontal (front to back) cooling
      - Several companies contacted about prototypes of vertically cooled shelves
    - Mechanical quality
  - AC/DC converters
  - IPMC
    - Agreed to test the IPMC designed at LAPP Annecy

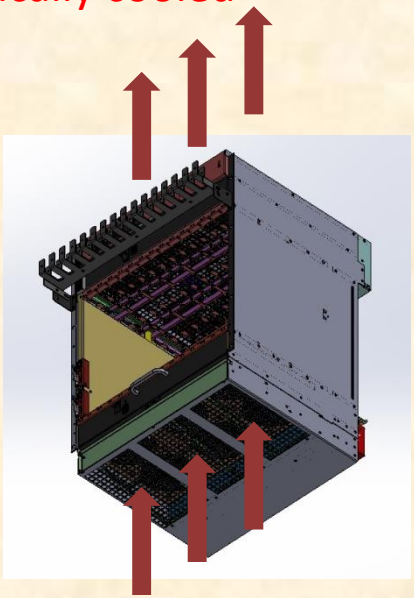

See also: https://indico.cern.ch/event/300897/contribution/7/material/slides/1.pdf

BE History & Plans

- One uTCA system purchased
  - Seen some interoperability problems
- Did not evaluate uTCA in detail as originally planned
  - For planning reasons only used uTCA as AMC backplane to host the GLIB card
  - May start evaluation this year!?

- A new Beam Position Monitor (based on the GLIB) for CLIC Test Facility is being designed in collaboration with LAPP

- Medium term plan:
  - Accelerator sector will put in place by mid-2014, in the context of the Controls Coordination Committee (CO3), a special working group for the recommendation of future Accelerator front-end platforms
  - Mandate and objectives not defined yet but the idea is to define the strategy for the next 15 years
  - Contribution form all CERN entities will be welcome. Various technologies will be considered, along with the associated operational deployment cost and migration paths

Contact: Marc Vanden Eynden and Frank Locci