advanced/micro TCA summit

Session 2-103

Sixty AdvancedTCA / MicroTCA Design Tips in 60 Minutes

Chuck Byers

Architect / Technical Leader

Cisco

Santa Clara, CA USA
November 1, 2011

advanced/micro TCA summit

Tip # 1:

Using Platforms

Use standards based platforms

- Open standards from organizations like PICMG or SAF form the best platform basis
- SCOPE can help select options
- Your goal should be 100% standards compliance…
- Except when you can't

advanced/micro TCA summit

Tip # 2:

Using Platforms

Use platforms to control total lifecycle cost of ownership

Without Common Platform
(**100 units** of total development cost)

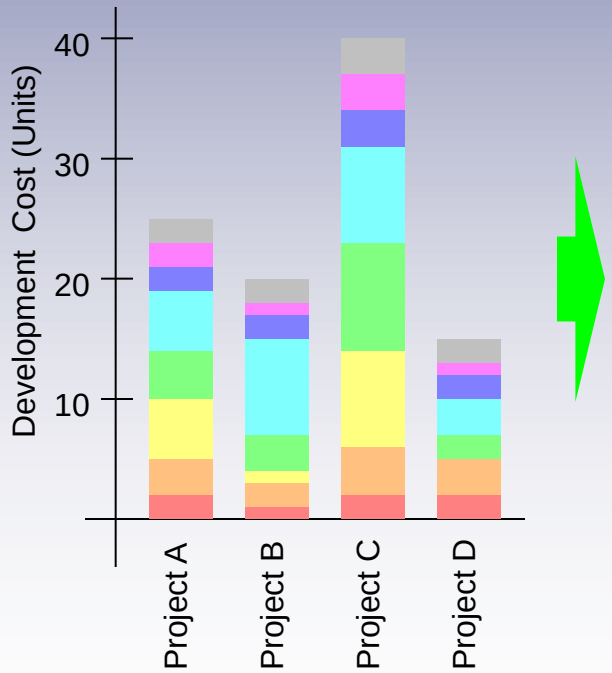

With Proprietary Platform
(**85 units** of total development cost)

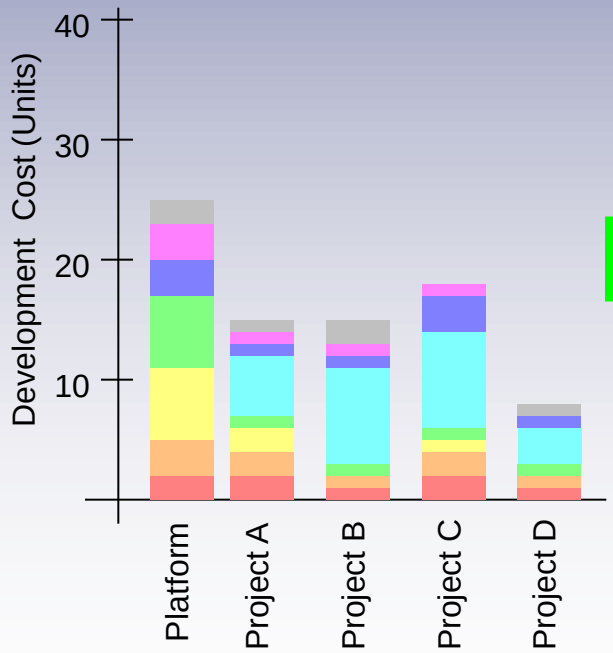

With Standard Platform
(**75 units** of total development cost)

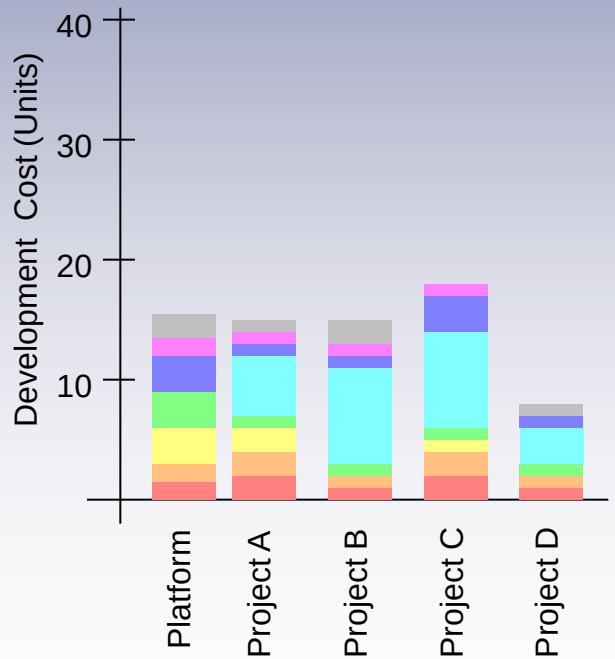

- Marketing
- Architecture
- HW Design
- Core SW
- App. SW
- Integration
- Factory
- Support

advanced/micro TCA summit

Tip # 3:

Using Platforms

Collect elements into a platform offer

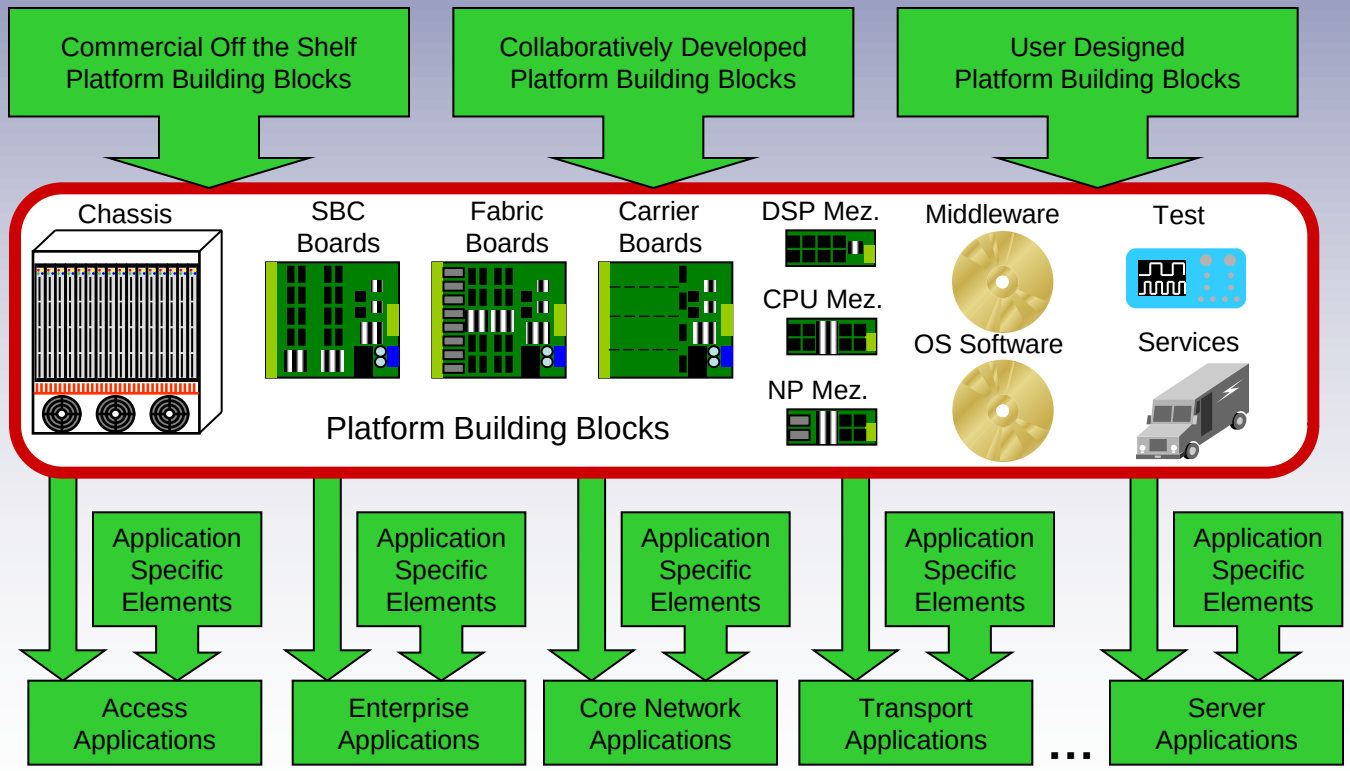

Image Courtesy of Alcatel-Lucent

advanced/micro TCA summit

Tip # 4:

AdvancedTCA or MicroTCA?

Choose AdvancedTCA or MicroTCA based upon the application

- AdvancedTCA is for larger scale systems with higher reliability or performance requirements
- MicroTCA is best for smaller scale or more cost critical applications
- The emerging AdvancedTCA Extensions specification may be optimal for enterprise and high density applications

advanced/micro TCA summit

Tip # 5:

AdvancedTCA or MicroTCA?

Consider derivative standards for special uses:

- AdvancedTCA Extensions for Enterprise uses
- AdvancedTCA for Physics
- Rugged MicroTCA – Air or Conduction Cooled
- Military versions of AdvancedTCA have been proposed

- There is also AXIe for Instrumentation and Test, but this is diverged from standard AdvancedTCA, and is being developed and supported outside PICMG.

advanced/micro TCA summit

Tip # 6:

COTS or Custom?

Use COTS HW and SW where feasible; custom where necessary

- COTS elements save development effort, improve schedule and reduce quality risk
- When you can't find a suitable product on the COTS market, you can design your own
- In-house designs can have better intellectual property control, and can keep trade secrets
- In high volume, in-house designs can cost less

advanced/micro TCA summit

Tip # 7:

COTS or Custom?

Partner with suppliers

- Some suppliers are open to untraditional supplier relationships, like contract design, royalties, or sliding price scales
- Suppliers are often looking for more of a partner relationship
- Often, it is desirable to have two suppliers for selected platform elements.

advanced/micro TCA summit

Tip # 8:

Partitioning

Identify natural modularity

- System partitioning should be driven by application partitioning
- Seek natural boundaries, and try to mimic them across the HW and SW
- Modularity can help optimally balance different types of system resources (CPU, DSP, NPU, Storage, I/O, etc.)

advanced/micro TCA summit

Tip # 9:

Partitioning

Note what types of processors your system needs

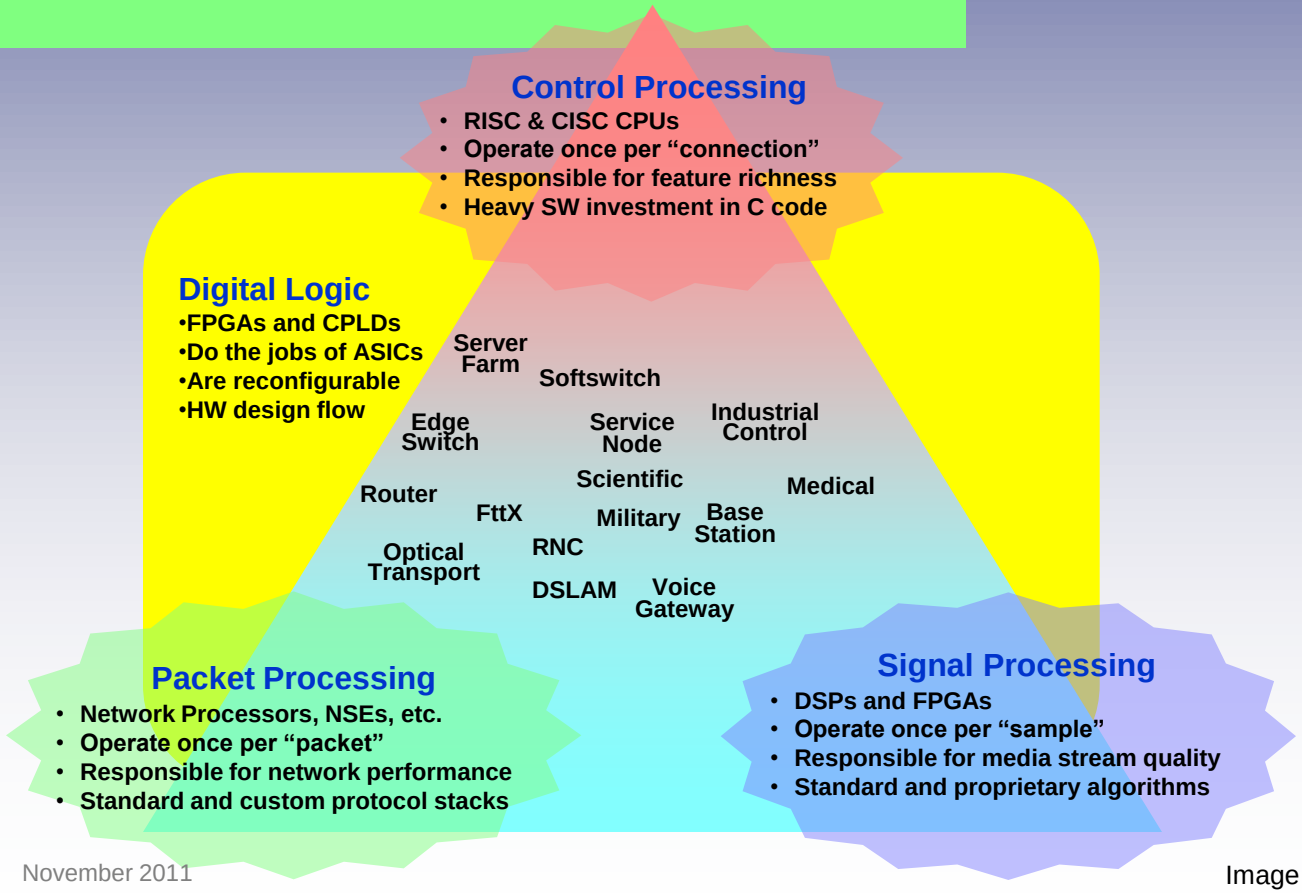

Image Courtesy of Alcatel-Lucent

advanced/micro TCA summit

Tip # 10:

Partitioning

Insure appropriate scalability

- Capacity
- Performance
- Reliability
- Feature richness
- All dimensions scale with cost points

advanced/micro TCA summit

Tip # 11:

Partitioning

Modularize Processors

- Processor technology is changing very rapidly
- Redesigning boards every ~12 months to accept the next processor technology is very expensive and risky
- Using smaller modules on a board can improve things
- AdvancedMCs are a reasonable choice for many CPUs
- COM Express modules are more optimized for processors
- For GPUs, consider MXM modules (commonly used to configure different levels of graphics in laptops)

advanced/micro TCA summit

Tip # 12:

Partitioning

Video requires lots of bandwidth per stream

Resolution	Horizontal	Vertical	Compressed Bandwidth (Mbps) H.264 AVC Videoconference	Uncompressed Bandwidth (Mbps) 8 bit color, 4:2:2 subsampling
QCIF	176	144	0.049	12
CIF	352	288	0.196	49
4CIF	704	576	0.782	195
WVGA	854	480	0.791	197
HD 720p	1280	720	1.778	442
HD 1080p	1920	1080	4.000	995

- Shelf I/O and backplane bandwidth must be adequate

advanced/micro TCA summit

Tip # 13:

Backplane Topology

Backplane topology follows application

- Systems with a tree approach to data processing are best served on dual star backplanes
- Systems characterized by lots of peer-to-peer traffic naturally fit full meshes or replicated meshes
- Systems with strong separation between control and data planes should consider dual-dual stars

advanced/micro TCA® summit

Tip # 14:

Backplane Topology

Enable mesh on backplanes

- The cost to add the traces and connectors to support a full mesh is often justified.
- This permits a 7.5X scaling in system capacity (three additional Moore's law cycles of system field life – 4.5-6 years)
- Modern switch chips make this cost effective
- May be able to use less risky link technology

advanced/micro TCA® summit

Tip # 15:

Subsidiary Specifications

Use PICMG 3.1 for all AdvancedTCA applications

- Ethernet has won the battle of the backplanes.
- 10G Ethernet is already cheaper per bit transported than GigE – prefer Option 9
- Soon, 10G Ethernet will be strictly less expensive to own than 1G interfaces (the way that GigE is cheaper than 100M Ethernet)
- 10G backplane links are necessary for many next generation services; 40G is coming soon

advanced/micro TCA® summit

Tip # 16:

Subsidiary Specifications

Don’t use obsolete specifications

- PICMG 3.3 and 3.6 are dead
- PICMG 3.2 and 3.5 are not well supported, but may be useful in special cases
- ATCA300 is D.O.A.
- Use caution with AMC.1 and AMC.4
- New versions of PICMG3.1 will support 40G

advanced/micro TCA® summit

Tip # 17:

Platform Elements

Build or buy AdvancedMCs to cover maximum applications

- Small and large CPU
- Network Processor
- Signal Processor
- Graphics / Video
- Disk
- Ethernet I/O

- These AdvancedMCs fit both in carrier boards and MicroTCA / PicoTCA systems

advanced/micro TCA summit

Tip # 18:

Platform Elements

Consider flash based storage

- Solid state disks offer advantages over rotating media like faster transfer speed, shock resistance, and longer MTBF
- However, their cost per GigaByte stored is ~8X more than conventional disks
- Single Level Cell (SLC) has better reliability, but higher cost than Multi-Level Cell (MLC)

advanced/micro TCA summit

Tip # 19:

Platform Elements

Hierarchical Storage

- RAM based caches
- Disks mounted on AdvancedTCA boards
- Disks on AdvancedMCs
- Disks on RTMs
- AdvancedMCs / RTM interfaces to JBODs
- Network Attached Storage

advanced/micro TCA summit

Tip # 20:

Platform Elements

Use network processors for packet interfaces

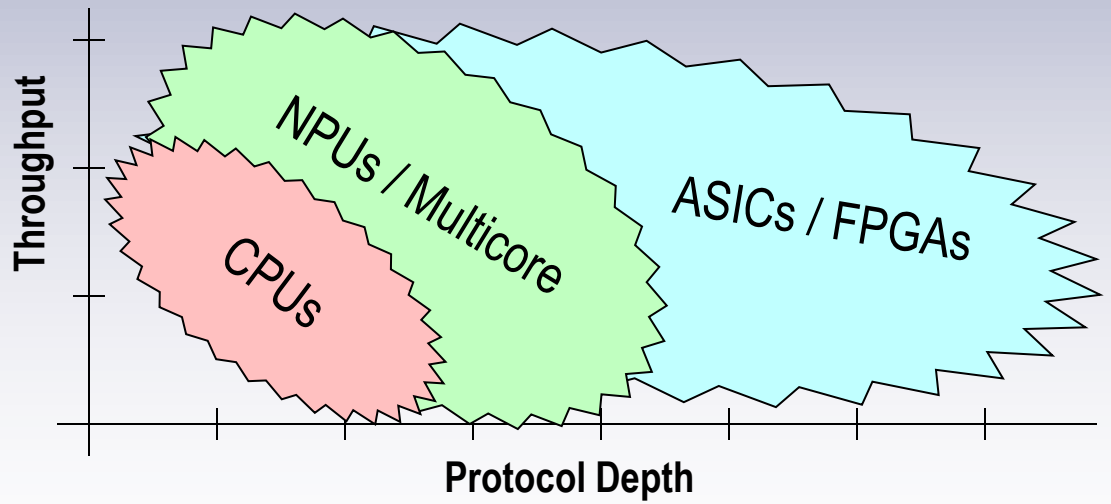

advanced/micro TCA summit

Tip # 21:

Platform Elements

Consider Multi-core Processors

- Today's IA processors with 4-6 cores
- PowerPCs with 2-8 cores, plus accelerators
- DSPs with 2-8 cores, plus special function units
- NPUs with up to 32 cores
- GPUs with up to 240 cores
- Parallel CPU chips with up to 100 cores

advanced/micro TCA summit

Tip # 22:

Mechanical

Vertical vs. horizontal packaging

- Vertical packaging has better airflow direction
- Vertical packaging has better cable management
- Horizontal packaging allows systems smaller than 8RU
- Horizontal packaging is more efficient for small systems

advanced/micro TCA® summit

Tip # 23:

Mechanical

Split a convertible AdvancedMC bay 6 ways

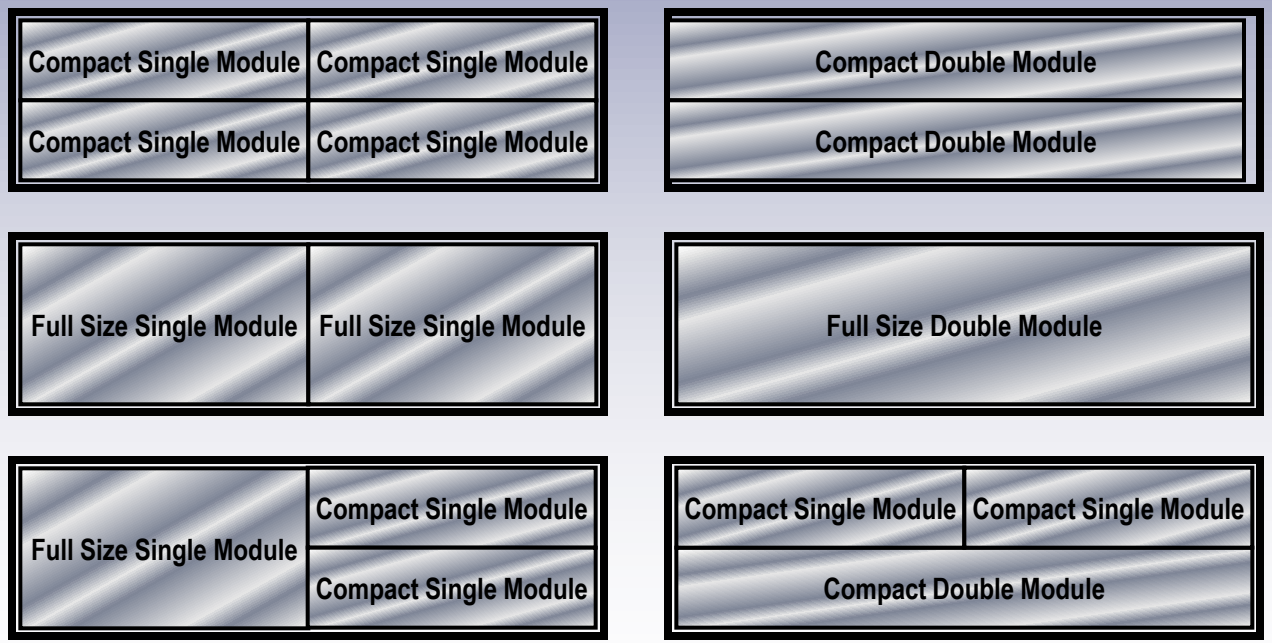

advanced/micro TCA summit

Tip # 24:

Mechanical

Consider double wide AdvancedTCA boards

- Huge capacity - Can hold:
  - Four 100W+ multicore CPUs with 64G+ RAM
  - 8 full size or 12 mid size AdvancedMCs
  - 72 RJ-45 jacks
  - 20TB of Disks (two layers of 5 * 2 TB drives)
- 600-800W Power dissipation
- Used by AdvancedTCA Extensions

advanced/micro TCA® summit

Tip # 25:

Mechanical

Cube packaging for MicroTCA

Simplex Cube Configuration

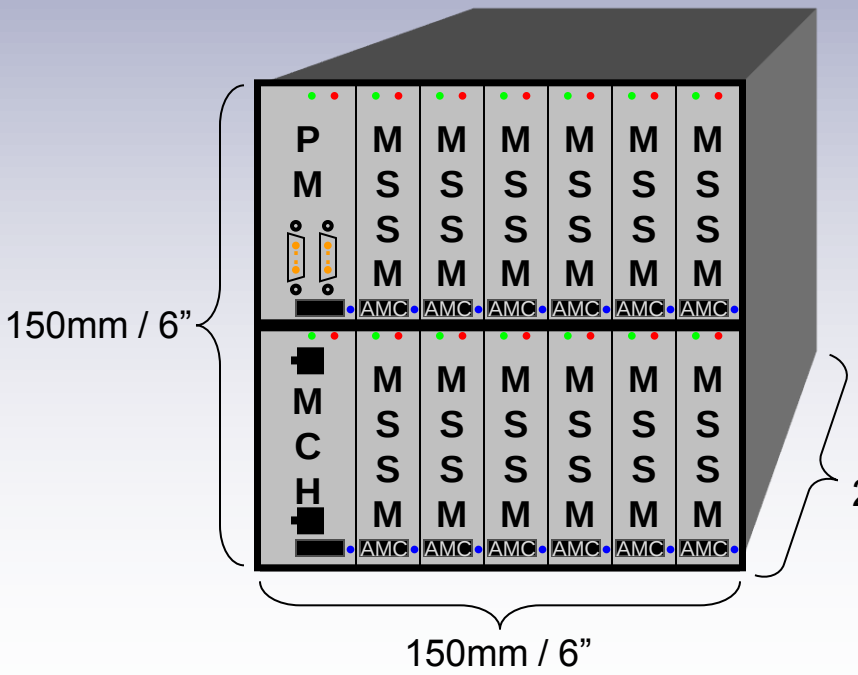

Duplicated Cube Configuration

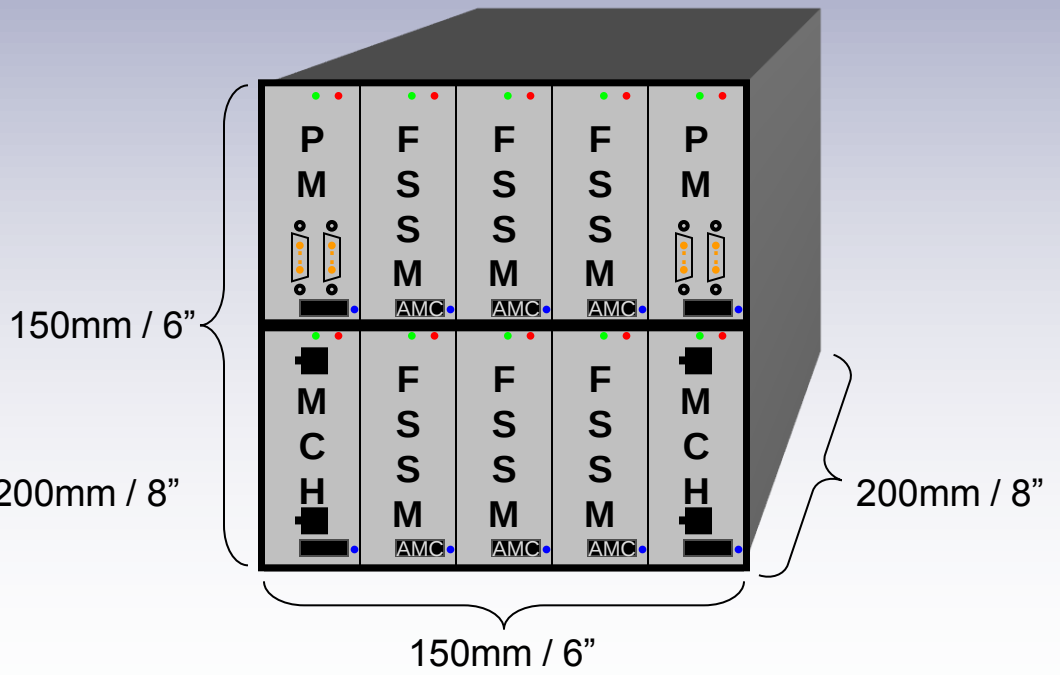

advanced/micro TCA summit

Tip # 26:

Mechanical

Cable management for 500+ fibers or CAT5 cables

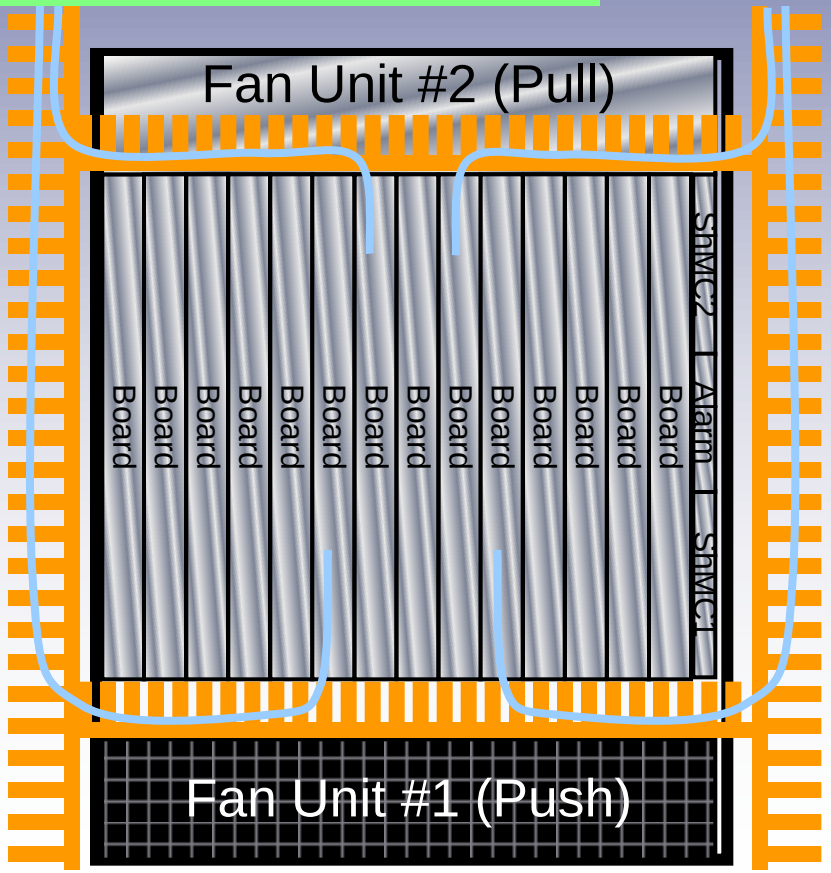

advanced/micro TCA summit

Tip # 27:

Mechanical

Extended Transition Modules

- The AdvancedTCA Extensions spec created something called an Extended Transition Module (ETM).
- An ETM connects to a front board via Zone 3, but is the same size as an AdvancedTCA board
- ETMs have extra board area and cooling to do many functions beyond the capabilities of RTMs:
  - Tens of gigabits of Deep Packet Inspection
  - 10+TB of disk
  - Hosting up to four full sized AdvancedMCs
  - Extra CPUs, GBUs, DSPs, memory, etc to supplement main board

advanced/micro TCA summit

Tip # 28:

Mechanical

Don’t forget special uses for Zone 3

- RTM / ETM interconnect
- Bulkhead connectors without RTM
- Supplementary backplane
- Optical backplane
- Additional power busses
- Cooling enhancements

advanced/micro TCA summit

Tip # 29:

Hardware Platform Management

Integrate ShMC onto fabric boards

- Shelf level ShMCs are a significant system cost contributor
- Their CPUs idle 99%+ of the time
- Fabric board processors have cycles to spare

advanced/micro TCA summit

Tip # 30:

Hardware Platform Management

Use HPM.1 for software update

- HPM.1 provides a standards based way to reliably and efficiently upgrade system software
- Also useful for FPGA configurations
- Don’t design systems where media like memory sticks or CDs must be physically transported to provide SW updates

advanced/micro TCA summit

Tip # 31:

Hardware Platform Management

Prepare for HPM.X Specs

- HPM.1: IPMI Controller Firmware Upgrade
- HPM.2: LAN-attached IPM Controller
- HPM.3: DHCP-assigned HPM Parameters
- HPM.4: Authenticating IPM Controllers

- These specifications are in various stages of completion by the HPM.X subcommittee
- Within ~year, most systems will need to comply

advanced/micro TCA summit

Tip # 32:

Power

Store holdup energy efficiently

- First, determine if holdup storage is even needed
- “Pump and dump” by charging HV capacitors to ~80V, and running power converters on this during dips
- Alternatively, use several Farads of ultracapacitor to store energy at ~3V

advanced/micro TCA® summit

Tip # 33:

Power

Do you really need 60V?

- Only a fraction of the world's COs (and no data centers) use 60V DC distribution
- More efficient and cheaper power trains can be designed if the maximum voltage is less than 75V
- The AdvancedTCA Extensions spec may eliminate the 60V supply option.

advanced/micro TCA® summit

Tip # 34:

Power

"Cheat" AdvancedMC power above 80W

- The 80W input power limit on AdvancedMCs limits their use (especially for PoE)
- Some models of AdvancedMC connectors (especially those used in MicroTCA) are capable of delivering more than 80W
- Larger power envelopes enable new apps.
- Better fans can cool >80W on bigger AMCs

advanced/micro TCA® summit

Tip # 35:

Power

Consider the power requirements of PoE

- Power over Ethernet (PoE) drives 48VDC down the same cable as the data interconnect
- Each 802.3af port supplies up to 15.4W (Future 802.3at standards extend this to above 50W)
- An AdvancedTCA board with 36 faceplate ports and 36 more on its RTM requires 1.1KW
- A full size AdvancedMC with 8 ports requires 123W+

advanced/micro TCA® summit

Tip # 36:

Power

Powering a ~15KW shelf

- Larger AdvancedTCA Extensions shelves could require almost 15KW (800W boards * 16 + two 1KW fan units)
- This is over 400A of -48V feed, requiring multiple circuits
- Running it on standard 120V circuits requires ~12 cords
- A better choice would be 208/240VAC, where about three 40A line connections would be required (2+1 redundant)
- An even better choice would be to consider the emerging 380VDC data center standard. This could be very efficient
- Next generation shelves should consider all these options

advanced/micro TCA® summit

Tip # 37:

Cooling

Use new generation of high flow / high pressure fans

- New fan lines from several manufacturers
- Move lots of air through high slot impedance
- Can enable 350W+ per slot on AdvancedTCA boards in non-NEBS environments
- Can enable ~80W on single wide AdvancedMCs
- Enables double wide AdvancedTCA Extensions boards of 800W+
- May improve shelf acoustic emissions

advanced/micro TCA summit

Tip # 38:

Cooling

Balance airflow between slots

- A system level problem
- Requires filler boards / AdvancedMCs in all open slots
- May require interchangeable restrictor plates
- Must address leaks around Zone 3 or RTMs
- Advanced systems may use variable dampers to exactly balance air between slots

advanced/micro TCA summit

Tip # 39:

Cooling

Design airflow at the system level

- Understanding the environment of a shelf is vital to successful cooling design
- Watch for frame-level recirculation
- High velocity hot exhaust streams can travel 10+ feet and enter intakes of other frames
- Side-to-side shelves suck in neighbors exhaust
- Understand wirecenter / datacenter cooling plan
- Energy savings by using more outside air for cooling

advanced/micro TCA® summit

Tip # 40:

Cooling

Advanced Cooling Techniques

- CP-TA classes B.1 through B.4 are inadequate for future systems. Work is underway to define classes C.5 through C.9+ in the AdvancedTCA Extensions subcommittee.
- Heat pipes and tall finned heatsinks can manage 100W+ chips (with adequate airflow)
- Liquid cooling is becoming an option for the most thermally challenged designs

advanced/micro TCA summit

Tip # 41:

Operating Systems

Use SAF and Linux

- Service Availability Forum middleware is highly optimized to work with AdvancedTCA and MicroTCA systems
- Linux is the de facto standard OS for nearly all xTCA applications
- Their performance, reliability and lifecycle cost have greatly improved in recent years

advanced/micro TCA summit

Tip # 42:

Operating Systems

Use virtualization

- Virtualization techniques can help balance the application's view of the system with the physical HW implementation
- Virtualization need not stop at traditional CPUs; it can potentially help manage pools of NPUs, DSPs and FPGA provided processing resources too, and also virtualize I/O resources

advanced/micro TCA summit

Tip # 43:

Application Software

XTCA in the Cloud

- Cloud services are not just about general purpose servers
- Packet processing, signal processing, storage, graphics, and parallel computation all have places in the cloud
- AdvancedTCA and MicroTCA are ideal modular platforms to provide highly configurable and reliable cloud services
- One design of chassis can host pretty much all the computational and storage resources needed for a cloud
- A chassis can be filled with multiple resource types, or several different types of shelf can host single resources

advanced/micro TCA summit

Tip # 44:

Application Software

Adapt 3rd party code

- Protocol stacks
- Drivers
- DSP algorithms
- Database management systems
- Security packages
- Desktop applications
- Web plug-ins

advanced/micro TCA summit

Tip # 45:

Application Software

Consider software reuse / multiuse

- Software at all levels of the stack should be designed for reuse and multi-use.
- When starting an application project, look around to identify existing SW that may work
- When coding new SW, design it in a structured, modular way that facilitates reuse
- Use good programming practice, especially documentation

advanced/micro TCA summit

Tip # 46:

Design Process

Design process rigor

- Use a rigorous, documented design process
- Establish entry / exit criteria for each phase
- Hold design reviews
- Plan for multi-group, multi-company, multi-continent, multi-timezone design teams
- Communication among all stakeholders is key!

advanced/micro TCA summit

Tip # 47:

Testing Techniques

Use JTAG for testing at all levels

- Stand alone AdvancedMC testing
- Testing full MicroTCA systems via JSM
- Testing collections of AdvancedMCs on carrier boards
- Testing full AdvancedTCA systems (requires change to the standard to replace test and ringing bus with JTAG interconnect, and probably the addition of a JSM)

advanced/micro TCA summit

Tip # 48:

Testing Techniques

Test and ringing buses are obsolete

- Since AdvancedTCA never got traction on metallic line units, the test and ringing busses in Zone1 are largely unused
- Don’t use these pins
- Future standards may reassign these pins to JTAG, USB, or other uses

advanced/micro TCA summit

Tip # 49:

Testing Techniques

Quality testing

- System quality can be tested in – to a point
- Bad architectures and bad implementations can’t be corrected by extensive test plans
- Testing only goes so far, then field experience takes over
- Get your system to field trial with a few friendly users ASAP.

advanced/micro TCA summit

Tip # 50:

Testing Techniques

Automate test suite

- Invest in test automation
- Regression test after every design change
- Where feasible, provide full load test suite
- Test development in parallel with code development is central to Agile SW methods
- Reliability testing (to a point)

advanced/micro TCA summit

Tip # 51:

Cost Reduction

Total Lifecycle Cost Of Ownership

- Total Lifecycle Cost of Ownership includes many components, all of which need optimization:
  - COGS
  - OPEX
  - Energy
  - Cooling
  - Installation
  - Updating
  - Growth / Degrowth
  - Maintenance
  - Repair
  - Evolution
  - Decommissioning
  - Recycling
  - Training

advanced/micro TCA summit

Tip # 52:

Cost Reduction

Use advanced cost reduction techniques

- Use most modern semiconductors (RAM, processors, flash, switches, power, etc.)
- Replace sheet metal with molded plastic (especially faceplates, and MicroTCA chassis)
- Consolidate fabric, ShMC
- Consider ASICs
- Reevaluate the need for redundancy

advanced/micro TCA summit

Tip # 53:

Time to Market

Get started with COTS

- Use fully-integrated COTS boxes for the most time critical applications – customize with SW
- Use COTS building blocks for prototypes
- As needed, transition selected elements from COTS to partner or in-house designs
- As volumes increase and decrease, COTS makes sense for different elements

advanced/micro TCA® summit

Tip # 54:

Future Evolution

Prepare for the new AdvancedTCA Extensions spec

Some concepts under consideration:

- 14-16 Double wide slots
- High performance cooling
- Up to 800W per board
- Improved power distribution
- Advanced RTM structures
- Platform management improvements

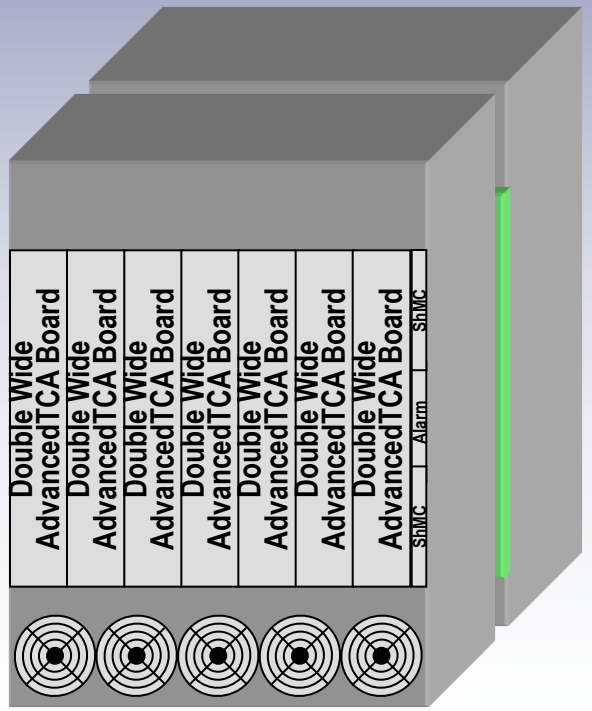

advanced/micro TCA summit

Tip # 55:

Future Evolution

Complete MicroTCA system on an AdvancedTCA Board
(Enabled by AdvancedTCA Extensions)

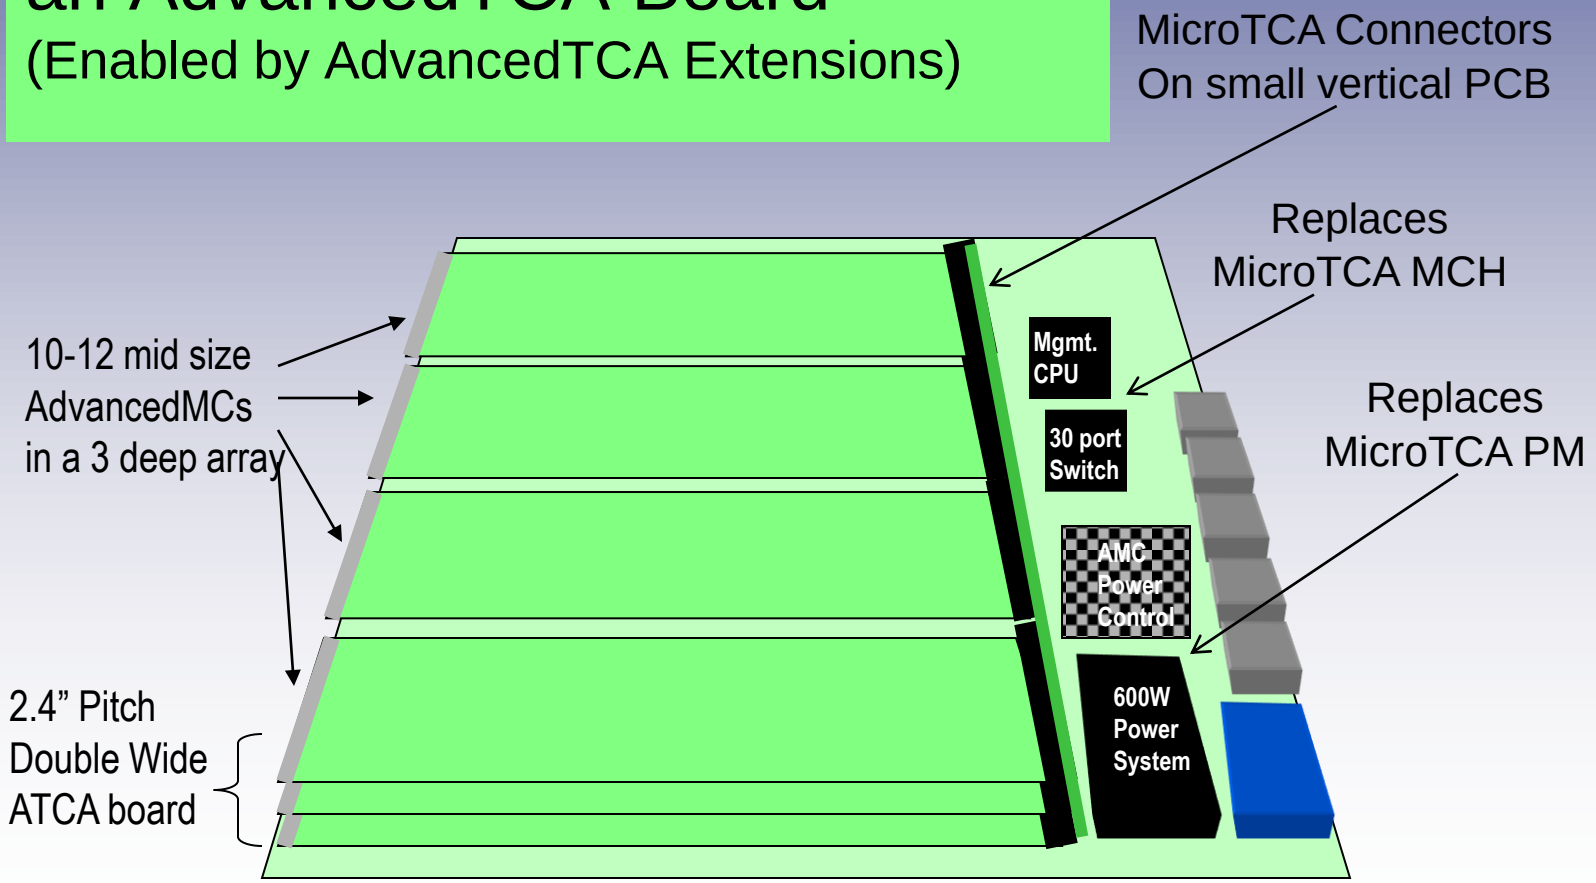

advanced/micro TCA summit

Tip # 56:

Future Evolution

Evolve toward 40G Backplane Interconnect

- 10Gb/s channels are common on backplanes, fabric boards and node boards (mostly XAUI)
- Modern switch chips will support 10GBaseKX4, permitting 40Gb/s on four lanes
- This increases a 14 slot AdvancedTCA main fabric bandwidth to 960Gb/s (dual star) or 7.28Tb/s (mesh)

advanced/micro TCA summit

Tip # 57:

Future Evolution

Supporting 100G Ethernet Ports

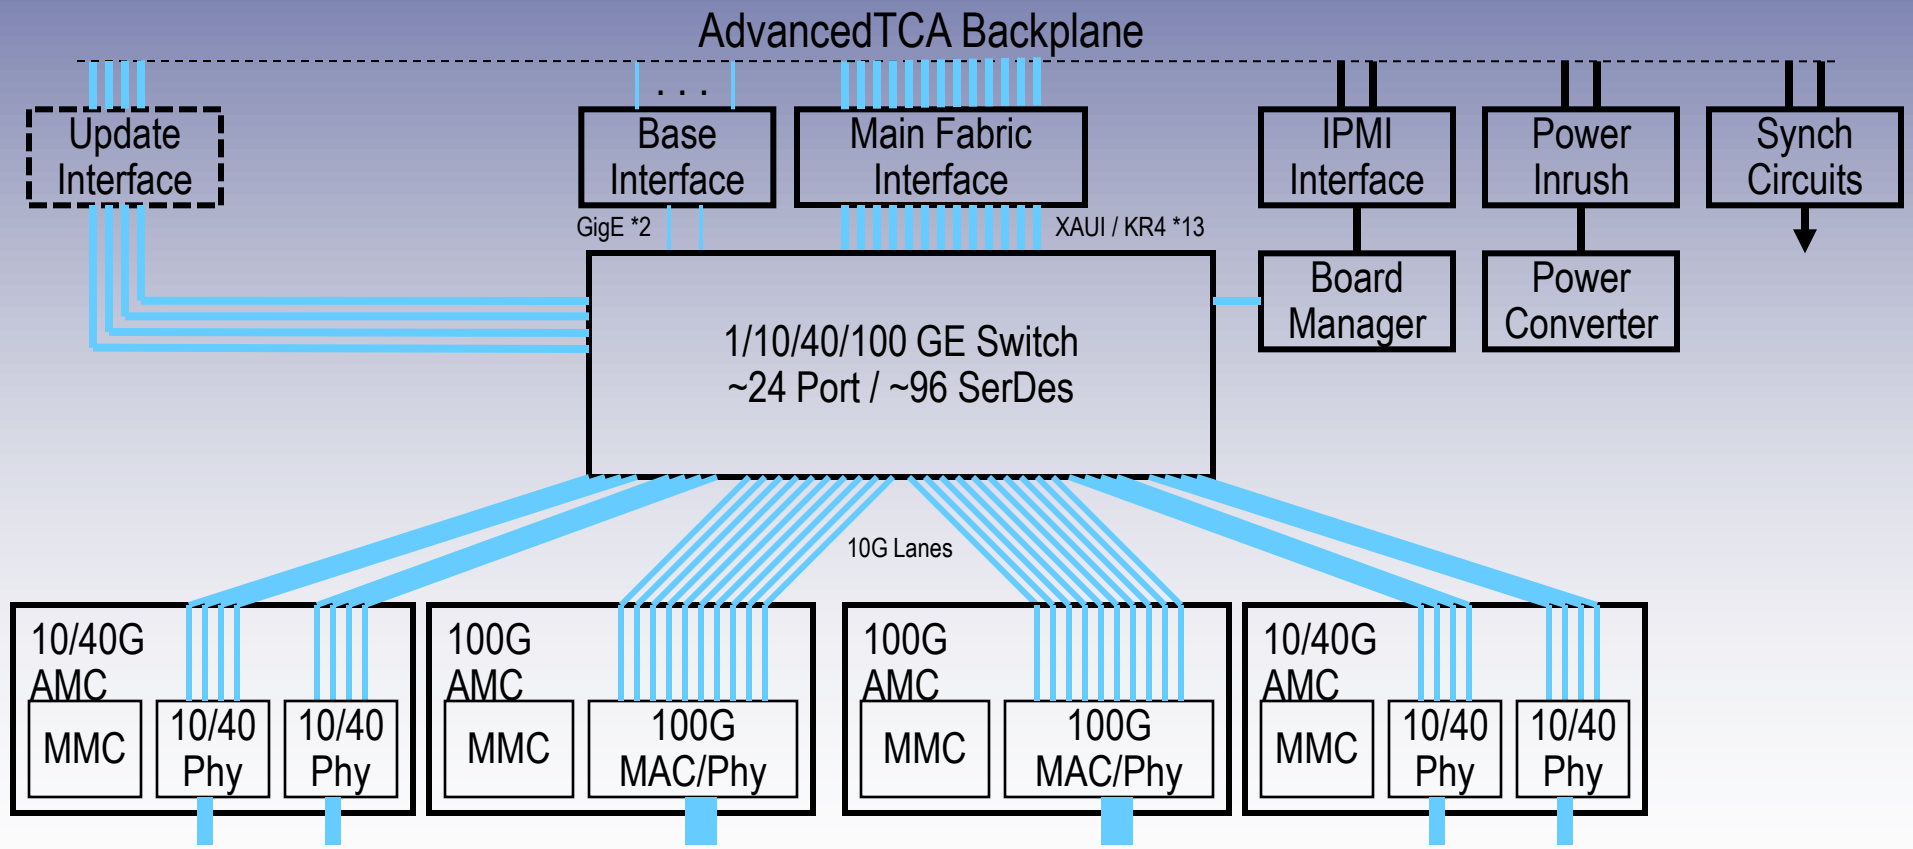

advanced/micro TCA summit

Tip # 58:

Future Evolution

Example Multi-core Processor Board

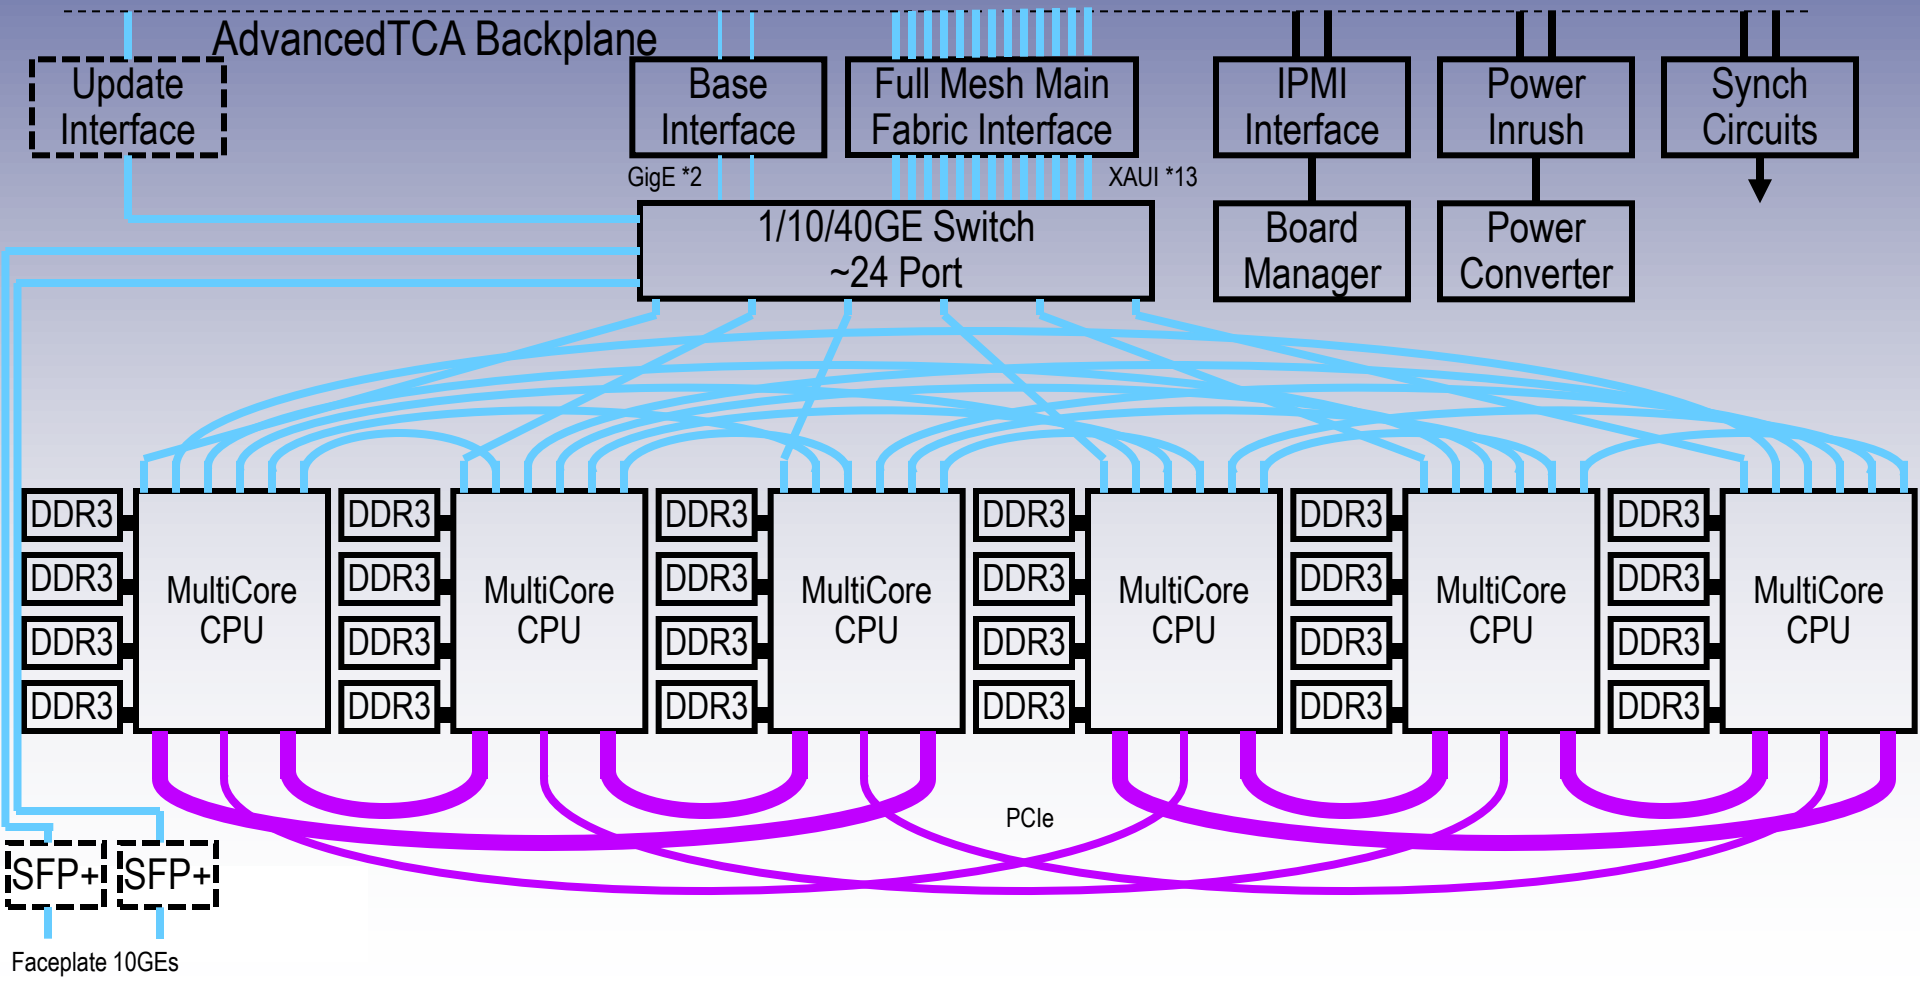

advanced/micro TCA summit

Tip # 59:

Future Evolution

Mesh of Meshes of Meshes

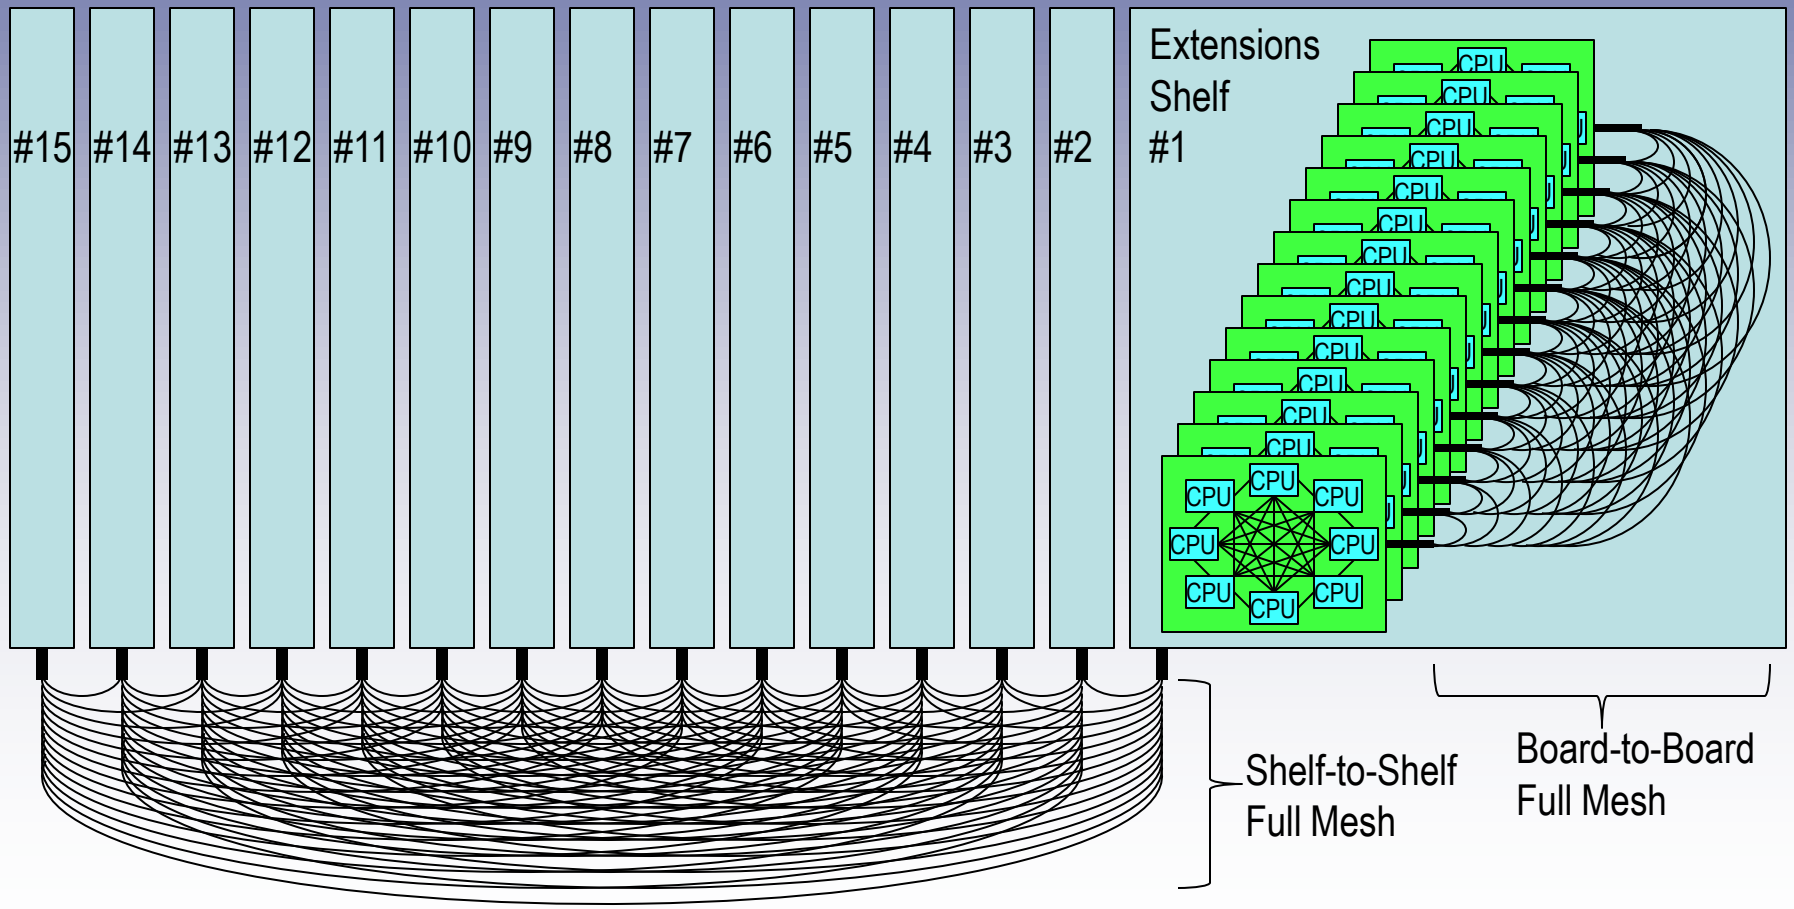

advanced/micro TCA summit

Tip # 60:

Future Evolution

Stick with the PICMG standards

- Over its 10 year history, adoption of AdvancedTCA hasn't been as fast as some have predicted
- Experience has shown that uptake is often slower than expected for new standards, but the total "area under the curve" often beats estimates
- As the global economic situation improves, spending on networks is going to accelerate, and AdvancedTCA / MicroTCA are ready to go.
- New specifications currently under construction will extend the life of PICMG standard based systems

advanced/micro TCA® summit

Thank You

Any Questions?

advanced/micro TCA® summit

Agenda:

1. **Using standards based modular platforms**
2. **AdvancedTCA vs. MicroTCA**
3. **COTS vs. custom design**
4. **System partitioning**
5. **Backplane topology choices**
6. **Which subsidiary spec to use**
7. **Elements (Fabrics, CPUs, DSPs, NPUs, I/O, Storage)**
8. **Mechanical packaging**
9. **Hardware platform management**
10. **Power subsystem design**
11. **Advanced cooling techniques**
12. **Middleware**
13. **Operating Systems**
14. **Applications Software**
15. **Design processes**
16. **Testing techniques**
17. **Cost reduction techniques**
18. **Time to market enhancement techniques**
19. **Future Evolution**
20. **Q&A**