Session 104

Sixty AdvancedTCA / MicroTCA Design Tips in 60 Minutes

Chuck Byers

Architect / Technical Leader

Cisco

advanced TCA summit

Tip # 1:

Using Platforms

Use standards based platforms

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

advanced TCA summit

Tip # 2:

CISCO

Using Platforms

Use platforms to control total lifecycle cost of ownership

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

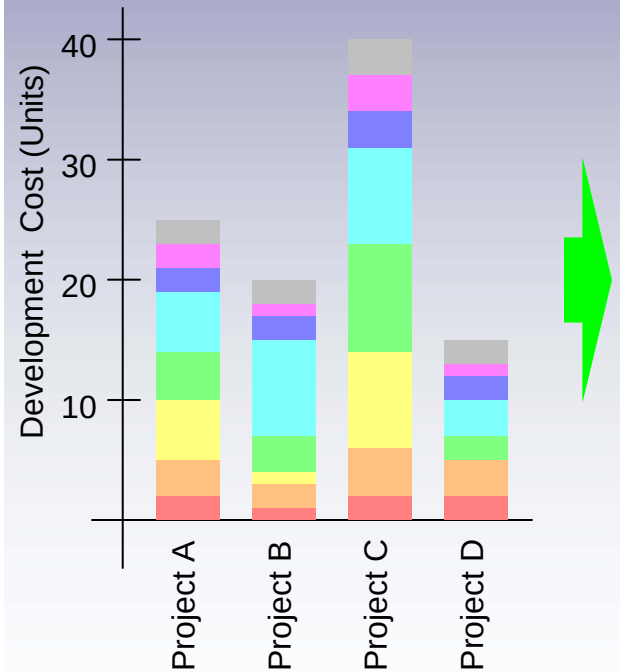

With Proprietary Platform
(**85 units** of total development cost)

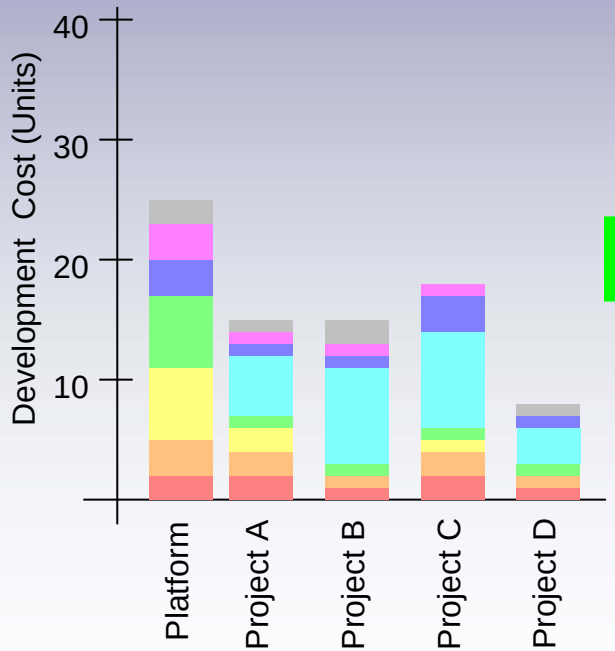

With Standard Platform
(**75 units** of total development cost)

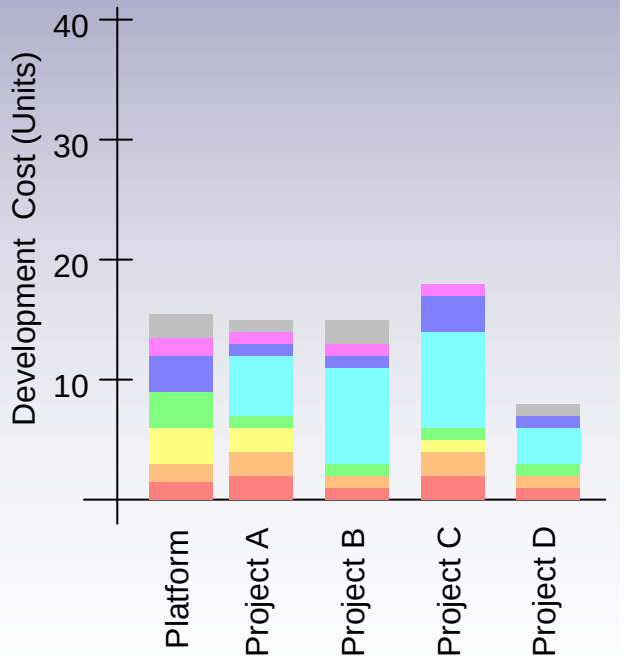

- Marketing
- Architecture
- HW Design
- Core SW
- App. SW
- Integration
- Factory
- Support

advanced TCA summit

Tip # 3:

CISCO

Using Platforms

Use the SCOPE Recommendations

- The PICMG standards are very versatile, but sometimes all those options complicate things and increase cost
- SCOPE has a series of recommendations and gap analysis documents that help define high volume options from the PICMG standards

advanced TCA summit

Tip # 4:

CISCO

Using Platforms

Collect elements into a platform offer

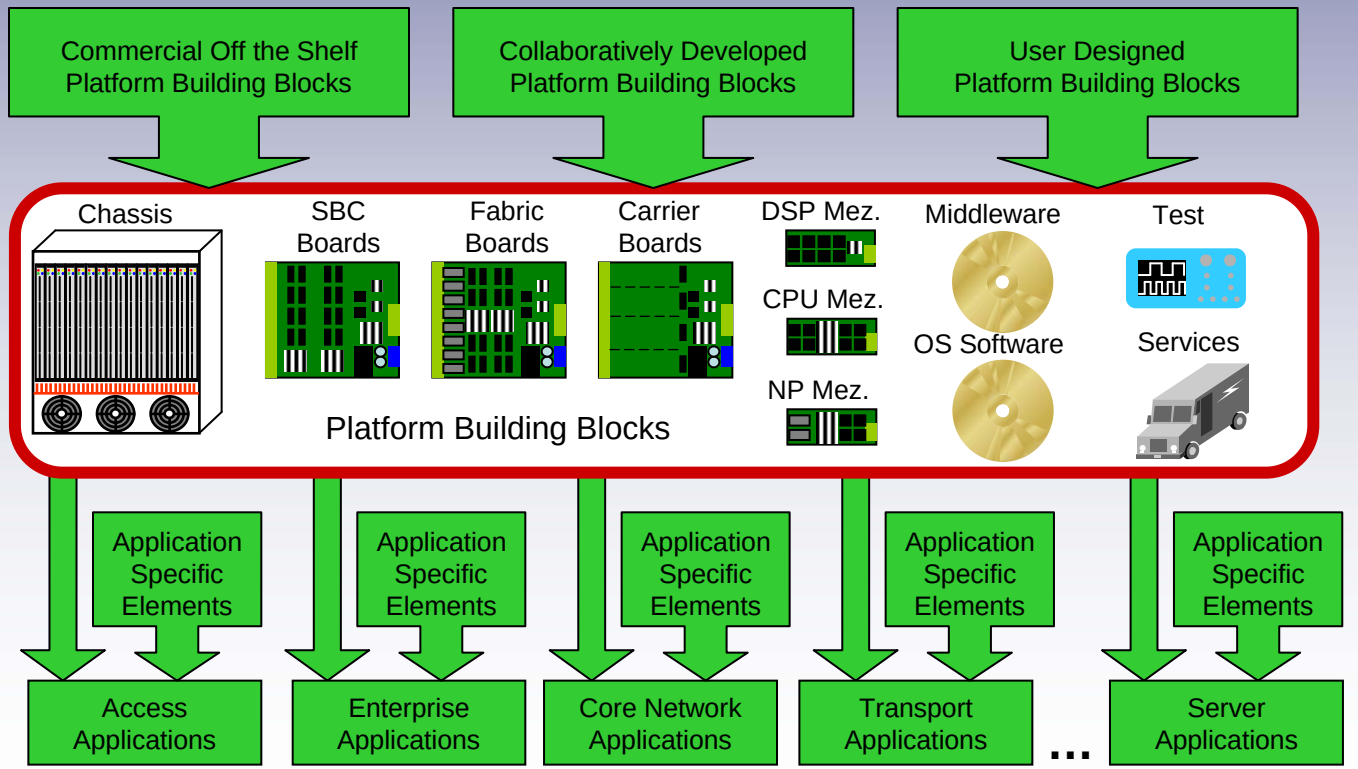

Image Courtesy of Alcatel-Lucent

advanced TCA summit

Tip # 5:

CISCO

Using Platforms

Same AdvancedMCs in multiple packaging options

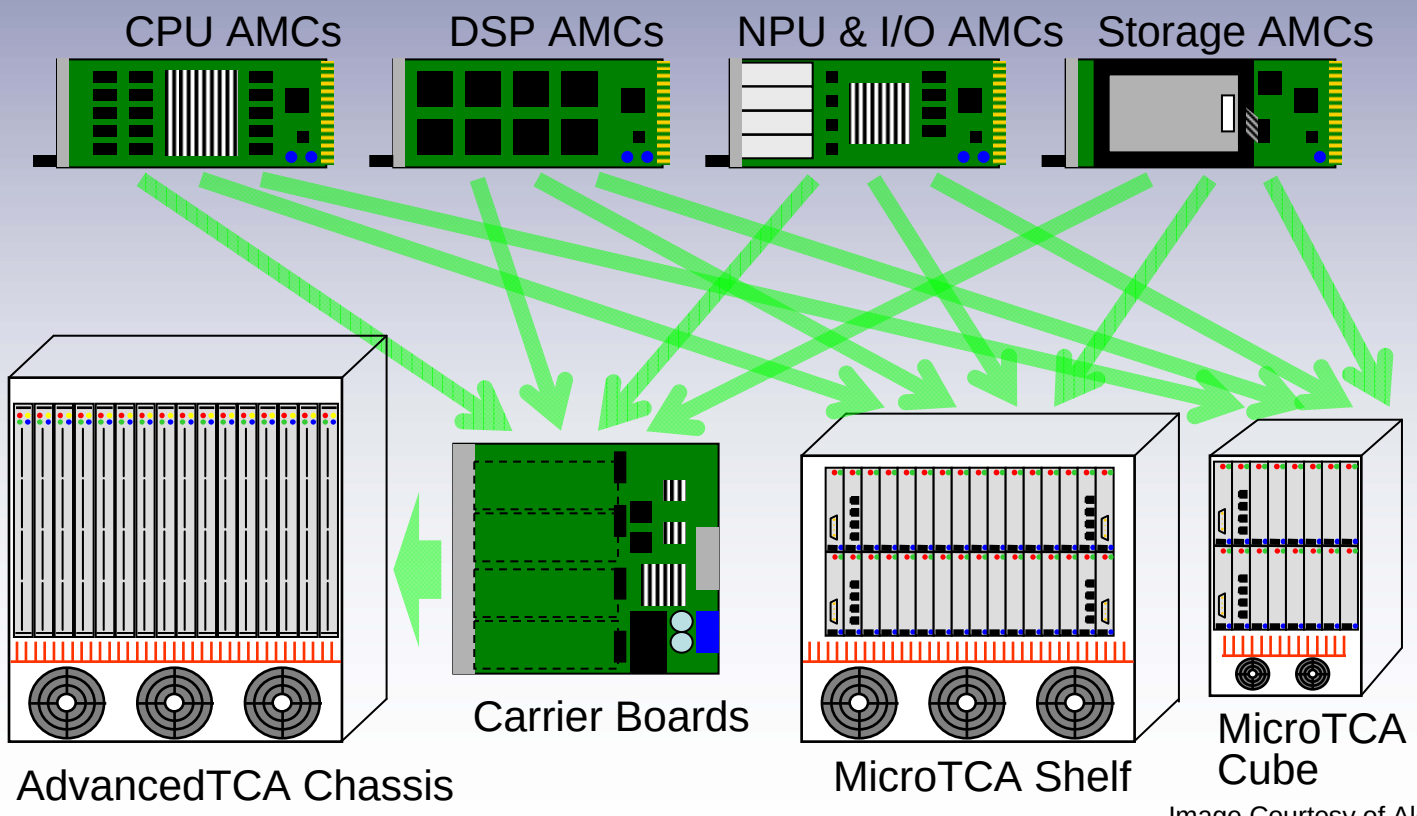

Image Courtesy of Alcatel-Lucent

advanced TCA summit

Tip # 6:

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

advanced TCA summit

Tip # 7:

AdvancedTCA or MicroTCA?

Use PicoTCA as a development tool

- Inexpensive to get started
- Every developer can have a small desktop system for development testing
- No contention for shared lab time

advanced TCA summit

Tip # 8:

CISCO

AdvancedTCA or MicroTCA?

“NanoTCA”

AdvancedMC Board

(CPU, NPU, DSP, FPGA, RAM, I/O, Storage, etc.)

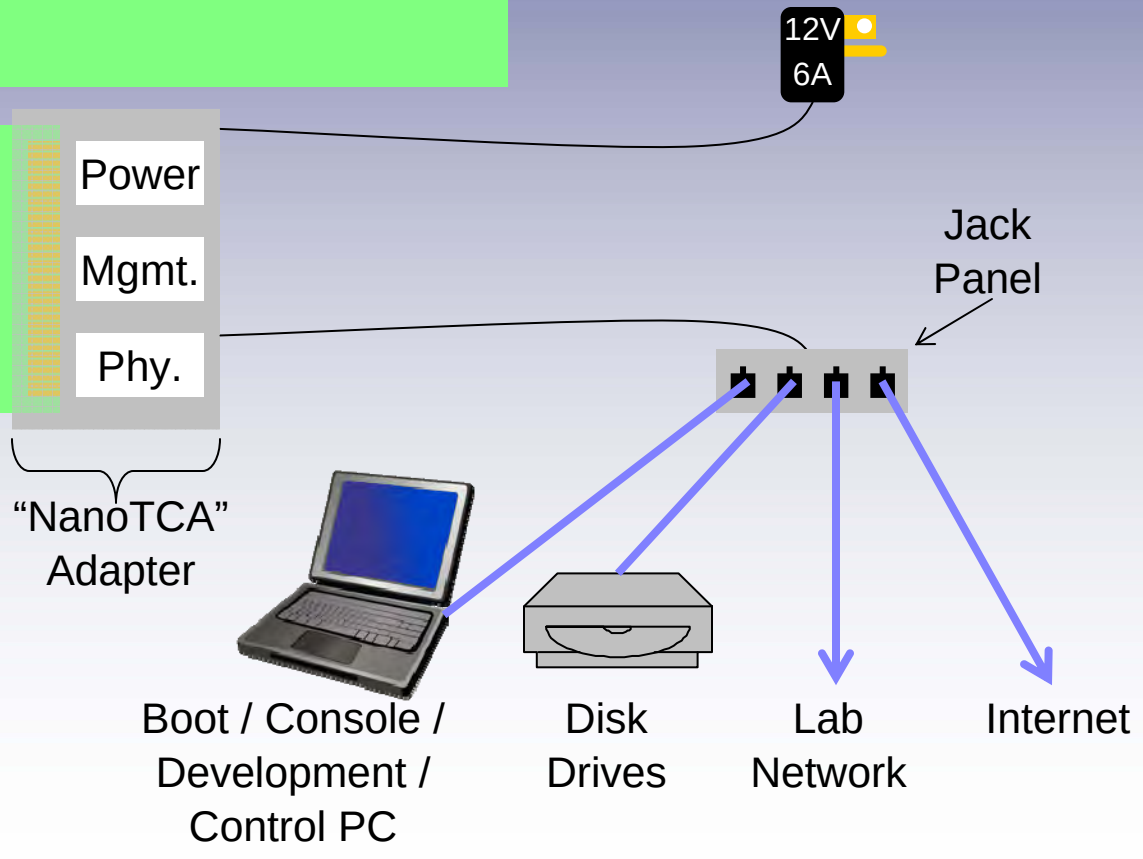

advanced TCA summit

Tip # 9:

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 TCA summit

Tip # 10:

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
- Partnerships strengthen supplier roadmaps by giving high visibility to future needs

advanced TCA summit

Tip # 11:

CISCO

COTS or Custom?

Make sure at least two COTS suppliers exist

- Each COTS element should have multiple sources of functionally equivalent products
- This protects you against supply chain, quality or performance problems
- Competition also keeps suppliers honest, by moderating their price and encouraging innovation

advanced TCA summit

Tip # 12:

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

advanced TCA summit

Tip # 13:

CISCO

Partitioning

Note what types of processors your system needs

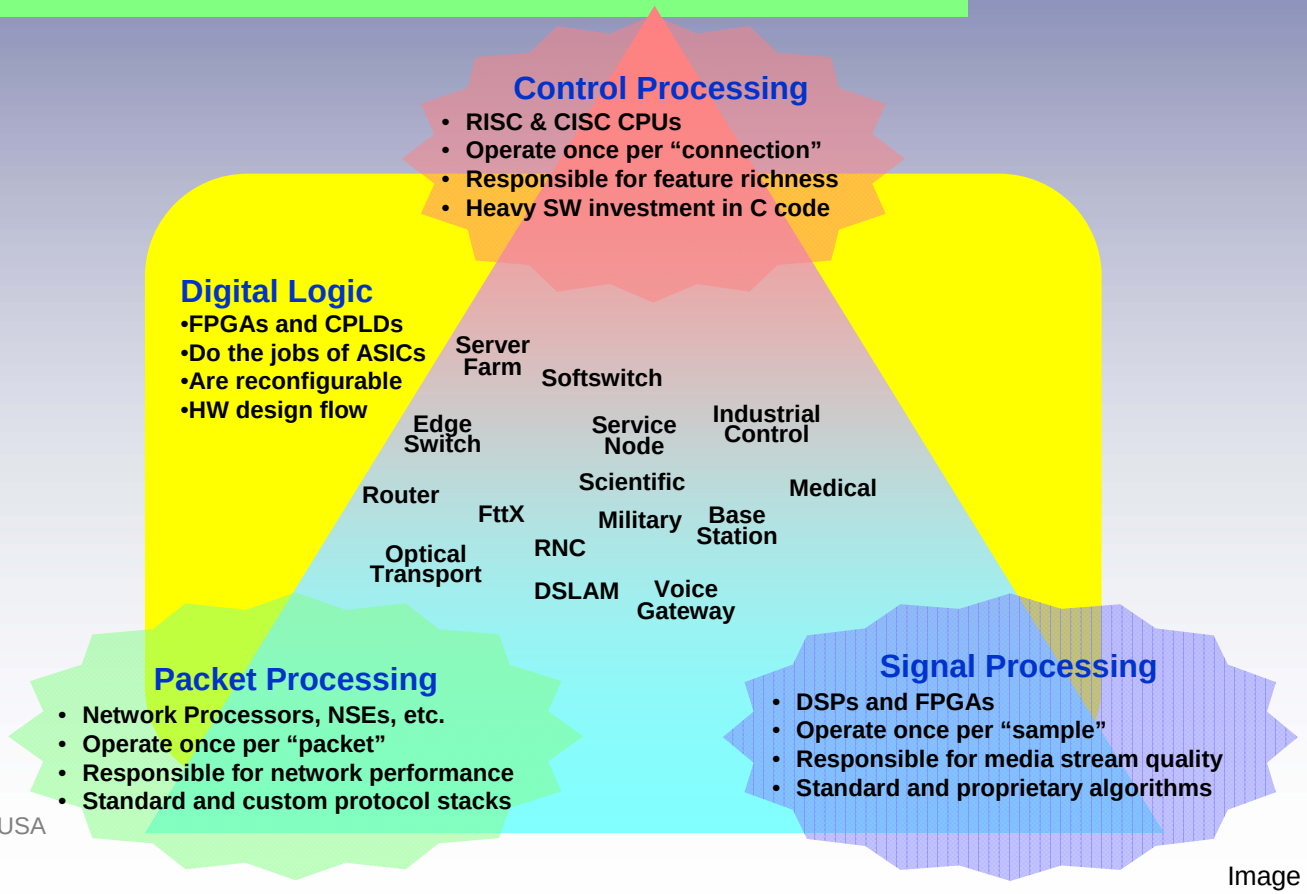

Image Courtesy of Alcatel-Lucent

advanced TCA summit

Tip # 14:

CISCO

Partitioning

Integrate multiple functions into single larger boxes

- AdvancedTCA / MicroTCA can host multiple boxes from the network block diagram in one shelf
- Functions can still be partitioned across boards
- Significant CAPX savings because elements like power cooling, and HPM aren't duplicated
- Significant OPEX savings

advanced TCA summit

Tip # 15:

Partitioning

Insure appropriate scalability

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

advanced TCA summit

Tip # 16:

CISCO

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
- Systems with strong separation between control and data planes should consider dual-dual stars

advanced TCA summit

Tip # 17:

Backplane Topology

Enable mesh on backplanes

- The extra 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)
- 2009 generation switch chips make this cost effective

advanced TCA summit

Tip # 18:

CISCO

Backplane Topology

Use replicated mesh in small shelves

- Replicated meshes can really increase performance in AdvancedTCA systems with 8 or fewer slots, at very little added cost
- Can eliminate the need for fabric boards
- Replicated mesh techniques may also apply to PicoTCA

advanced TCA summit

Tip # 19:

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

advanced TCA summit

Tip # 20:

CISCO

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 probably D.O.A.
- Use caution with AMC.1 and AMC.4

advanced TCA summit

Tip # 21:

Platform Elements

Build or buy AdvancedMCs to cover maximum applications

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

advanced TCA summit

Tip # 22:

CISCO

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 20-30X more than conventional disks
- Single Level Cell (SLC) preferrable to Multi-Level Cell (MLC) due to reliability advantages

advanced TCA summit

Tip # 23:

CISCO

Platform Elements

Multiple alternatives for hosting disk drives

- Directly mounted on AdvancedTCA boards
- On AdvancedMCs
- On RTMs
- Small flash drives may even fit on MCHs and ShMCs
- In specialized MicroTCA storage cubes

advanced TCA summit

Tip # 24:

CISCO

Platform Elements

Use network processors for packet interfaces

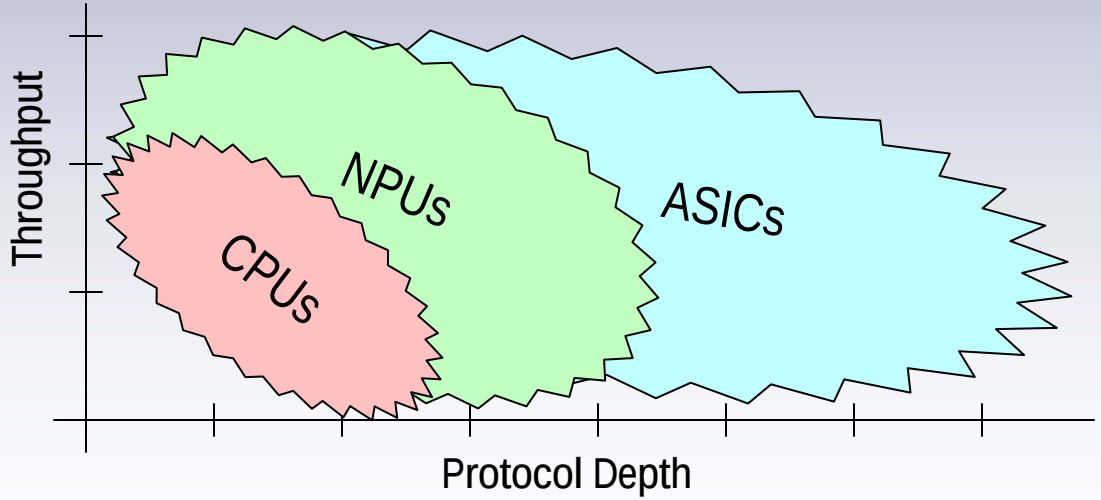

advanced TCA summit

Tip # 25:

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 TCA summit

Tip # 26:

CISCO

Mechanical

Split a convertible AdvancedMC bay 6 ways

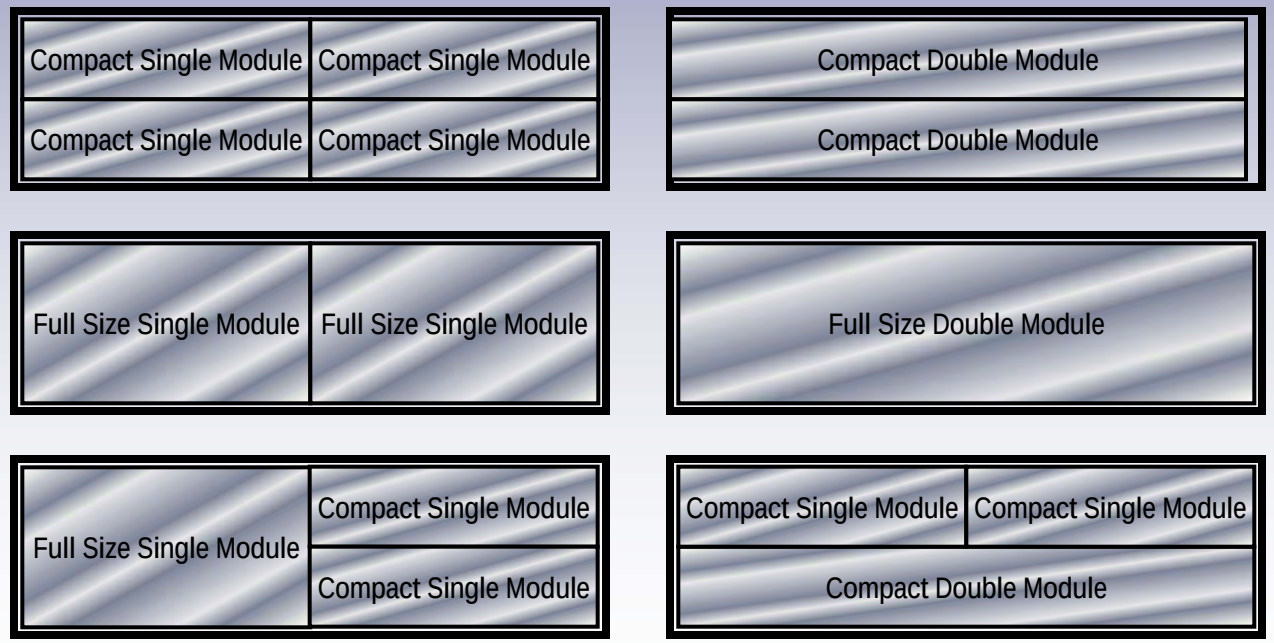

advanced TCA summit

Tip # 27:

CISCO

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
  - 10TB of Disks (two layers of 5 * TB drives)
- 600W Power dissipation

advanced TCA summit

Tip # 28:

CISCO

Mechanical

Cube packaging for MicroTCA

Simplex Cube Configuration

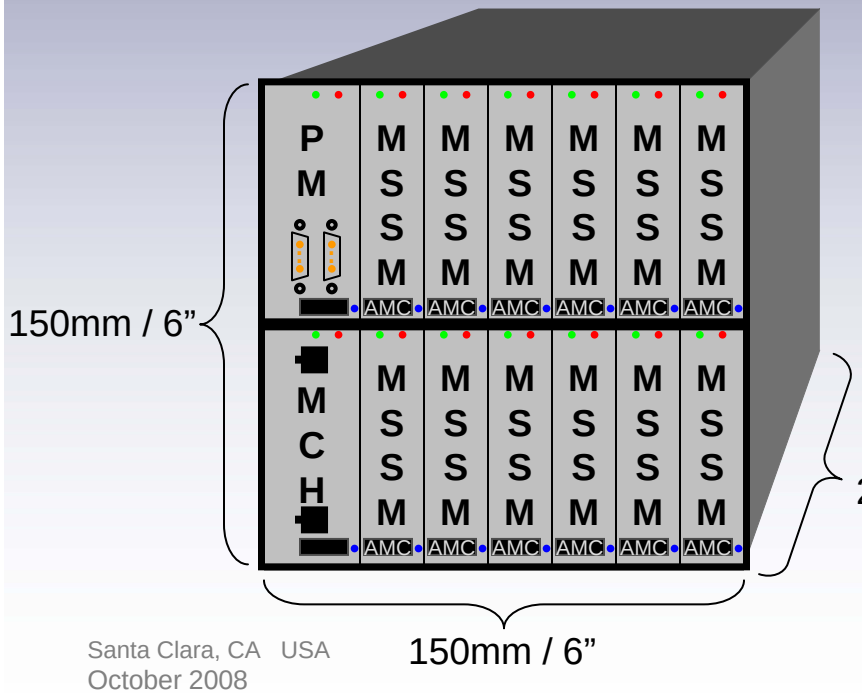

Duplicated Cube Configuration

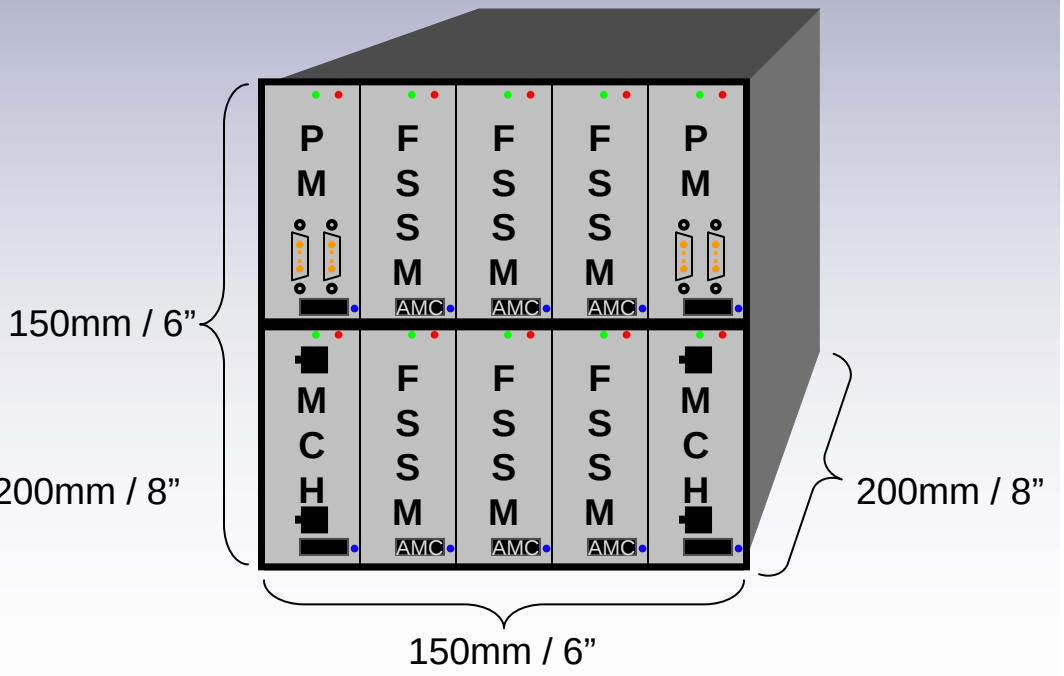

advanced TCA summit

Tip # 29:

CISCO

Mechanical

Cable management for
500+ fibers or CAT5 cables

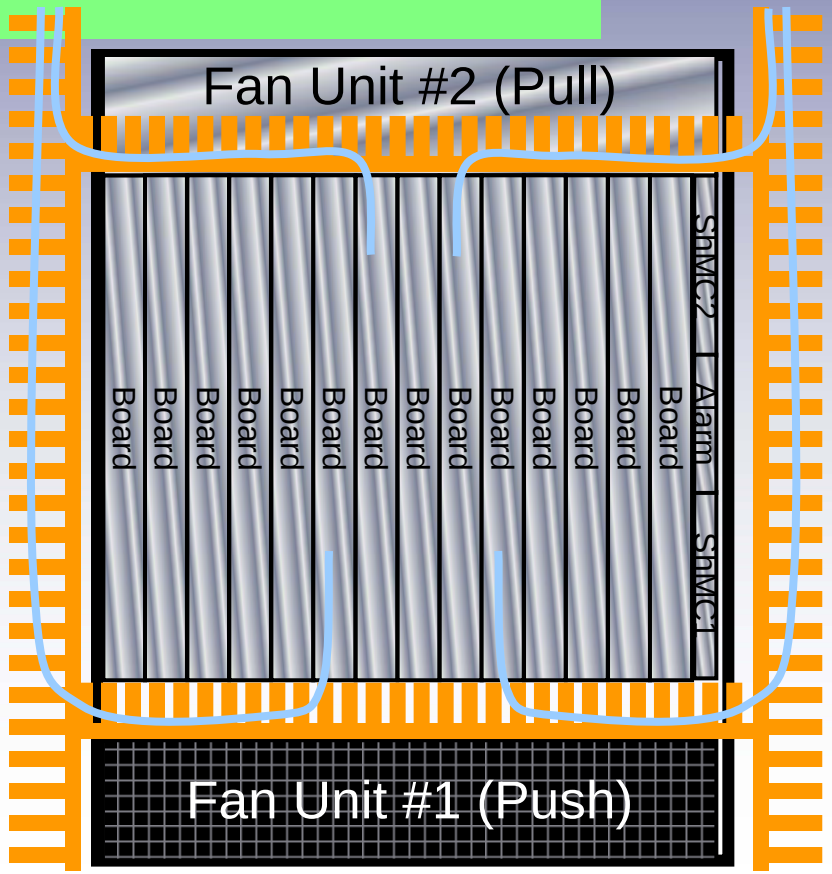

advanced TCA summit

Tip # 30:

Mechanical

Don’t forget special uses for Zone 3

- RTM interconnect
- Bulkhead connectors without RTM
- Supplementary backplane
- Optical backplane
- Additional power busses

advanced TCA summit

Tip # 31:

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 TCA summit

Tip # 32:

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 TCA summit

Tip # 33:

Hardware Platform Management

Supplement IPMI links

- Don’t try to transport high bandwidth traffic over the system’s IPMI links
- Use base interface / common options for higher bandwidth management traffic
- Reserve IPMI links for high priority and latency sensitive control messages

advanced TCA summit

Tip # 34:

CISCO

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 TCA summit

Tip # 35:

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

advanced TCA summit

Tip # 36:

CISCO

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 TCA summit

Tip # 37:

Power

Consider the power requirements of PoE

- Power over Ethernet (PoE) drives 48VDC down the same cable as the data interconnect
- Each port supplies up to 15.4W (Future PoE standards may extend this to about 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 TCA summit

Tip # 38:

CISCO

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 AdvancedTCA boards in non-NEBS environments
- Can enable ~80W on single wide AdvancedMCs

advanced TCA summit

Tip # 39:

CISCO

Cooling

Balance airflow between slots

- A system level problem
- Requires filler boards / AdvancedMCs in all open slots
- May require interchangeable restrictor plates
- Advanced systems may use variable dampers to exactly balance air between slots

advanced TCA summit

Tip # 40:

CISCO

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

advanced TCA summit

Tip # 41:

Middleware

Use standards based middleware

- Service Availability Forum middleware is highly optimized to work with AdvancedTCA and MicroTCA systems
- Several suppliers with compliant offers
- OpenSAF is an emerging standard for open source middleware conforming to the SAF standard – Encourage its uptake (Linux model)

advanced TCA summit

Tip # 42:

Middleware

Use middleware to control configuration / clones

- Middleware can control which boards / modules will operate in a system
- Use this to limit configurations to those that have been tested
- Also permits box level suppliers to limit boards / modules that work in a box to only their own
- If advanced crypto is used, very tough to break

advanced TCA summit

Tip # 43:

Operating Systems

Use Linux

- Linux is the de facto standard OS for nearly all xTCA applications
- Its performance and reliability have greatly improved in recent years

advanced TCA summit

Tip # 44:

CISCO

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

advanced TCA summit

Tip # 45:

Application Software

Adapt 3rd party code

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

advanced TCA summit

Tip # 46:

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 TCA summit

Tip # 47:

CISCO

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

advanced TCA summit

Tip # 48:

CISCO

Design Process

Close process feedback loops

- Encourage bidirectional communication between all process levels:

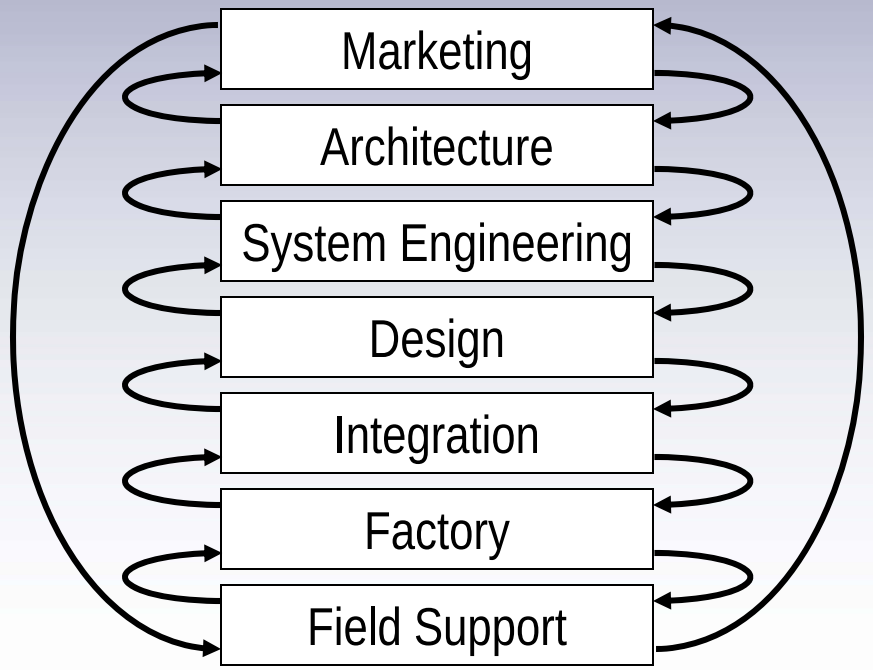

advanced TCA summit

Tip # 49:

CISCO

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 TCA summit

Tip # 50:

CISCO

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 or other uses

advanced TCA summit

Tip # 51:

CISCO

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

advanced TCA summit

Tip # 52:

CISCO

Testing Techniques

Automate test suite

- Invest in test automation
- Regression test after every design change
- Where feasible, provide full load test suite
- Reliability testing (to a point)

advanced TCA summit

Tip # 53:

CISCO

Cost Reduction

Continuously evaluate COTS strategy

- As volumes increase and decrease, COTS makes sense for different elements
- As suppliers enter and leave competitive markets, cost profiles change

advanced TCA summit

Tip # 54:

CISCO

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 TCA summit

Tip # 55:

CISCO

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

advanced TCA summit

Tip # 56:

Time to Market

Compromise when necessary

- Often, compromises are necessary to hit a fixed market window. Dimensions can include:
  - Reliability
  - Feature richness
  - Ease of use
  - Capacity
  - Performance
  - Cost

advanced TCA summit

Tip # 57:

CISCO

Future Evolution

Complete MicroTCA system on an AdvancedTCA Board

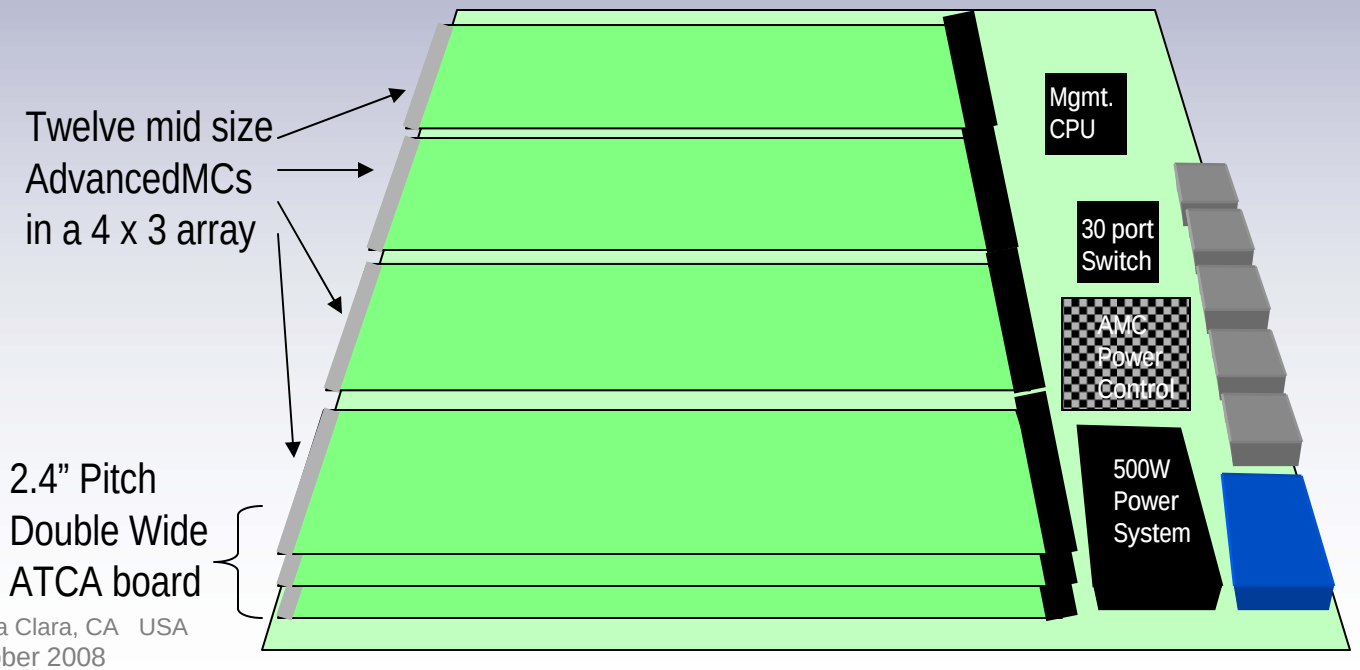

advanced TCA summit

Tip # 58:

CISCO

Future Evolution

Evolve to back-to-back AdvancedTCA

- 14-16 Double wide slots
- Push/pull cooling
- 600W per board
- No RTMs or Zone 3
- ShMCs in Front
- PEMs in rear (4 * 50A)

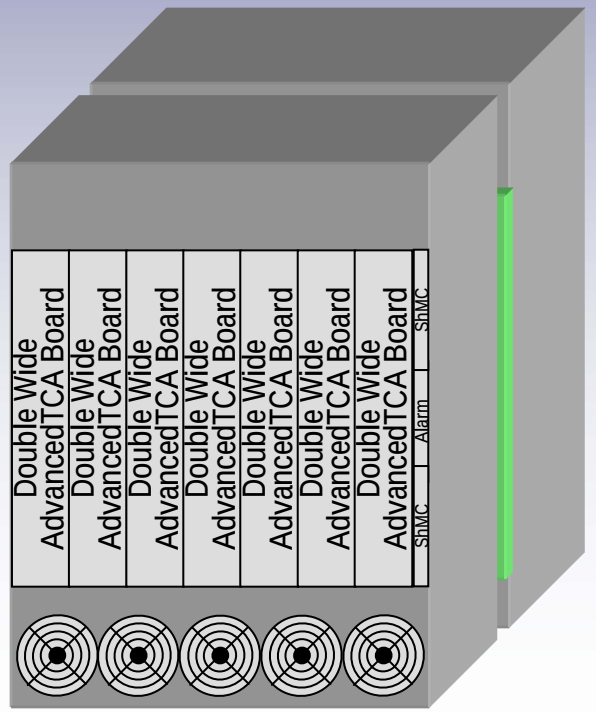

advanced TCA summit

Tip # 59:

Future Evolution

Consider fluid cooling

- Forced air is “out of gas” for the most demanding systems
- Fluid cooling systems have thousands of times the heat removal capacity of air
- Radiators can be placed in various places in the system to move heat between air and fluid
- Some hot chips can be directly fluid cooled

advanced TCA summit

Tip # 60:

Future Evolution

Stick with the PICMG standards

- Adoption of AdvancedTCA and MicroTCA hasn't been as fast as some have predicted
- History has shown that uptake is often slower than expected for new standards, but the total "area under the curve" often beats estimates
- Once the global economic situation improves, spending on networks is going to accelerate, and AdvancedTCA / MicroTCA is ready to go.

Thank You

Any Questions?

Santa Clara, CA USA
October 2008

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