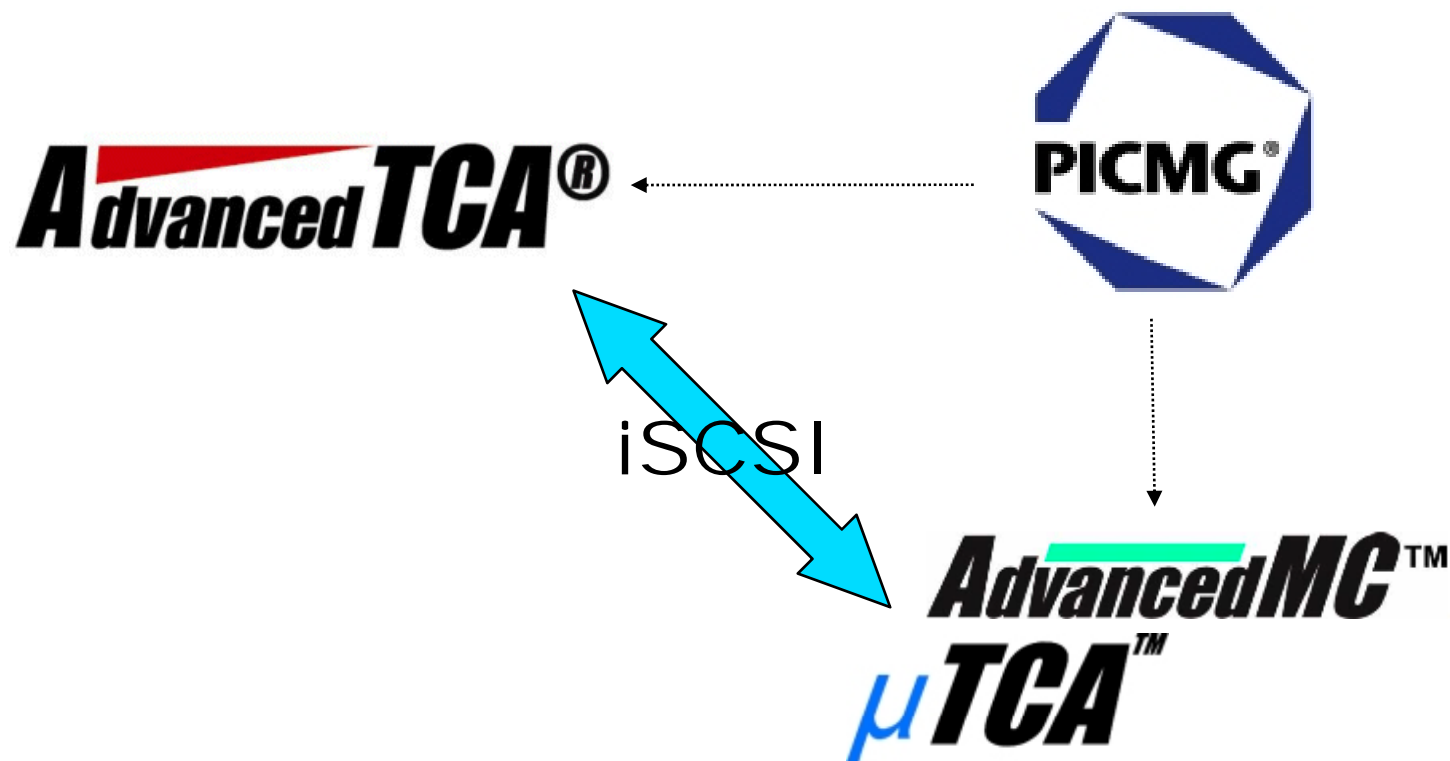


iSCSI Storage for AdvancedTCA



Introduction

- Why Storage?
- Storage Choices
- Proposal: iSCSI
- iSCSI Solutions

• Why storage?

- AdvancedTCA offers a broad and capable architecture for packet-transfer system builders
- System vendors want to deliver higher value functions as well within the ATCA ecosystem
 - Internet Media Servers (IMS)
 - WiMax Gateways with caching and co-located service delivery
 - Dense cluster computing to compete with BladeCenter[™] and other vendor specific formats
- To enable these and other applications the system must provide storage in a way that complements the architecture

•Storage Choices

- Local storage, on-blade/on-AMC
 - Consumes area, power and complexity
 - Limited number of disks
- In-Chassis storage
 - Fabric compatibility
 - Bandwidth mismatch
- Off-Chassis storage
 - Enclosure type mismatch
 - Lack of standards

• Local Storage

- Many dedicated processing ATCA blades offer on-blade storage
 - If the disks are mounted directly on the board they become non-FRU
 - If mounted in an AMC bay they consume payload slots better used by I/O
- Typical ATCA blade has 2-4 AMC bays
 - Limit of 2 disks - Storage redundancy options limited to RAID-1 mirroring
 - Limit of 4 disks -RAID-5/6 is now possible
- Few ATCA processing blades can fit 4 AMC payload bays, a Raid-On-Chip HBA and significant CPU resources
 - Must consume CPU resources for RAID functions
 - Software complexity

• In-Chassis Storage

- Fabric compatibility issues
 - Fabric option specifications are mutually exclusive
 - FibreChannel, PCIe, 10Gb XAUI and SRIO all use the same fabric lanes by default
 - Mixing protocols would require non-standard lane selections
 - FibreChannel may be a good storage fit, but it fails to provide a good match to packet switching
 - 10Gb XAUI is gaining steam, matches the packet switched world and supports iSCSI
- Bandwidth mismatch

•Off-Chassis Storage

- Enclosure mismatch – ATCA to ???
 - JBODs and RBOD enclosures are not PICMG and generally consumer or enterprise class (not NEBS!)
 - Not designed for FRU/front-panel service
- Lack of standards
 - Every hard disk enclosure company has their own formfactor, sled, cabling, controller CPU blade etc.
 - Not plug-and-play, not interchangeable

• Proposal: iSCSI


- IETF Standard
- Native on 10Gbps XAUI ethernet
- Supported by major operating systems (Windows, Linux, Solaris etc.)
- Comparable market in the enterprise arena offers life-cycle support from many vendors and critical-mass for the technology

• iSCSI Solutions

- In-Chassis iSCSI server
 - Bandwidth mismatch
- Off-Chassis iSCSI server
 - How to get out of the box?
 - Will enterprise class solutions be acceptable?

• Bandwidth Mismatch

- ATCA provides as much as 160Gbits/second or more of fabric interconnect within a chassis
- A single node blade within an ATCA chassis is typically limited to a single fabric port from the fabric switch blade
 - Ex: 10Gb XAUI – 13 blades @ 10Gbs feeding one storage blade at 10Gbs – 12 blades are starved for storage or every blade gets $\frac{3}{4}$ Gbits/second



• How to get out of the box?

- Use NIC adapters for each blade
 - A rats nest of front panel cabling
- Use a storage-enabled fabric switch
 - On-board RAID controller?
 - Not available on the market
- Use fabric blade up-link ports to cross connect to another fabric enabled chassis

• A PICMG Solution

- MicroTCA shares the ancestry and much of the specification support from ATCA
 - Preserves the investment in IPMI and other systems tools and knowledge
 - Preserves FRU/hot-swap functionality

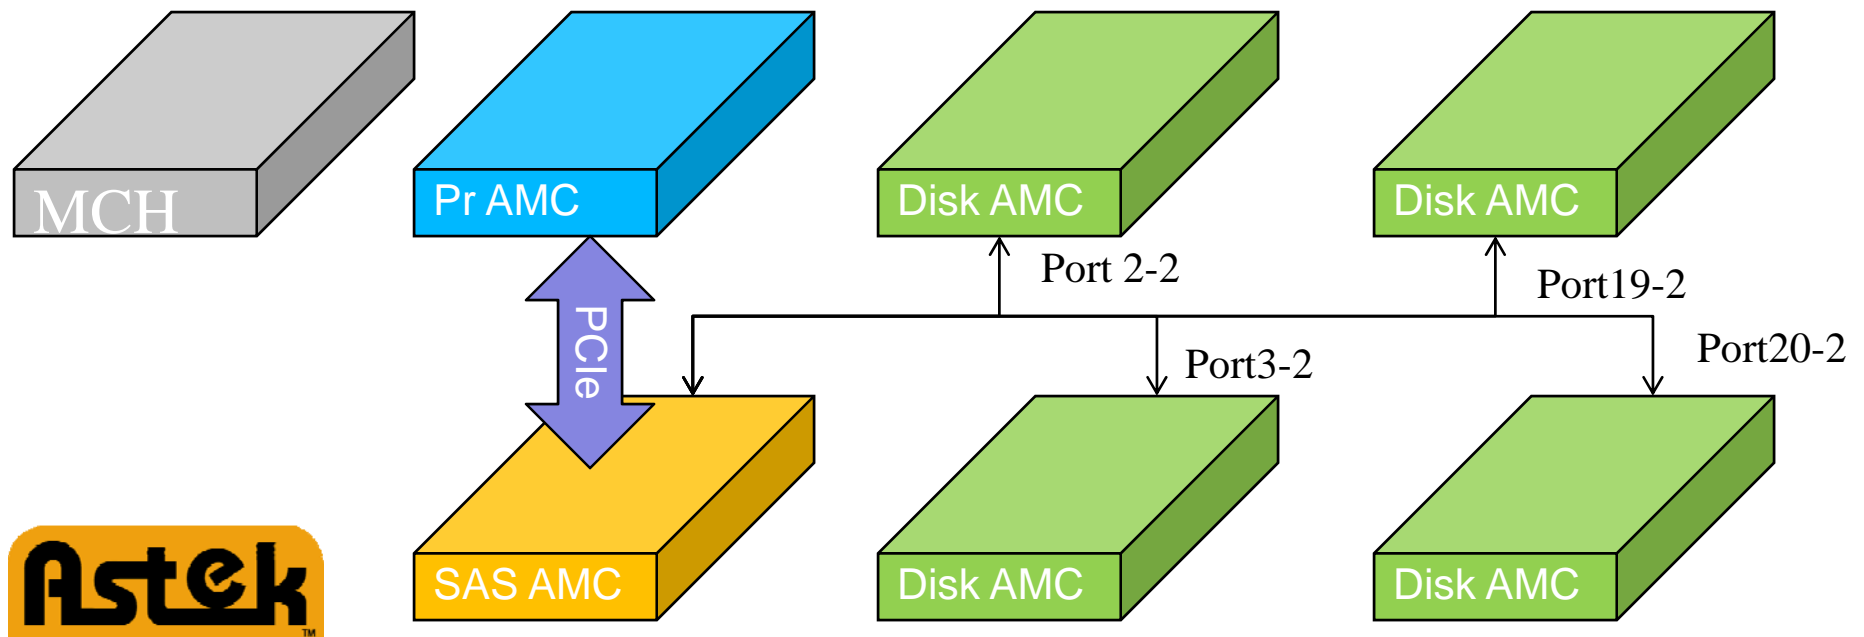
• Build-Up

- MicroTCA chassis with fully routed Ports2/3 to left and right MCHs
- MCH cards with XAUI uplink and a SAS expander on-board
- Processor AMC(s) with XAUI and PCIe fabric inputs
- SAS initiator AMC
- Up to 10 disk AMC cards

• Example System



- MicroBlade/ELMA MicroBox 1U chassis
- Power, MCH, PrAMC, SAS AMC, 4 Disk AMCs
- iSCSI target running on PrAMC
- 4 disks + external JBOD if needed



•What's Next

- Storage enabled MCHs
- Fully routed MicroTCA backplanes
- Combined AMC with XAUI NIC, processor and SAS initiator
- Mate storage intelligence solutions (SAN managers) with IPMI intelligence (SAF implementations)

• Conclusion

- Storage is critical to new platform applications
- Most current choices don't scale well
- iSCSI fits well
- Solutions can be built today
- Vendors can create new solutions

- Questions?
- Copies of this presentation, and a white paper detailing the assembly of a MicroTCA storage server at:
<http://www.astekcorp.com/>

