

LAN-attached Management Controllers: Challenges, Solutions, and a PICMG Update

Tutorial T1A – New Directions in xTCA™

November, 2010

Mark Overgaard



World-Class Management Components

FOCUSED. DEPENDABLE. PROVEN.

- Introduction to Pigeon Point Systems
- LAN-attached management controller candidates in xTCA
- Motivations and various implementation approaches for LAN attach connections
- Update on LAN-attach-related activities of the PICMG Hardware Platform Management Extensions (HPM.x) subcommittee
- Conclusion

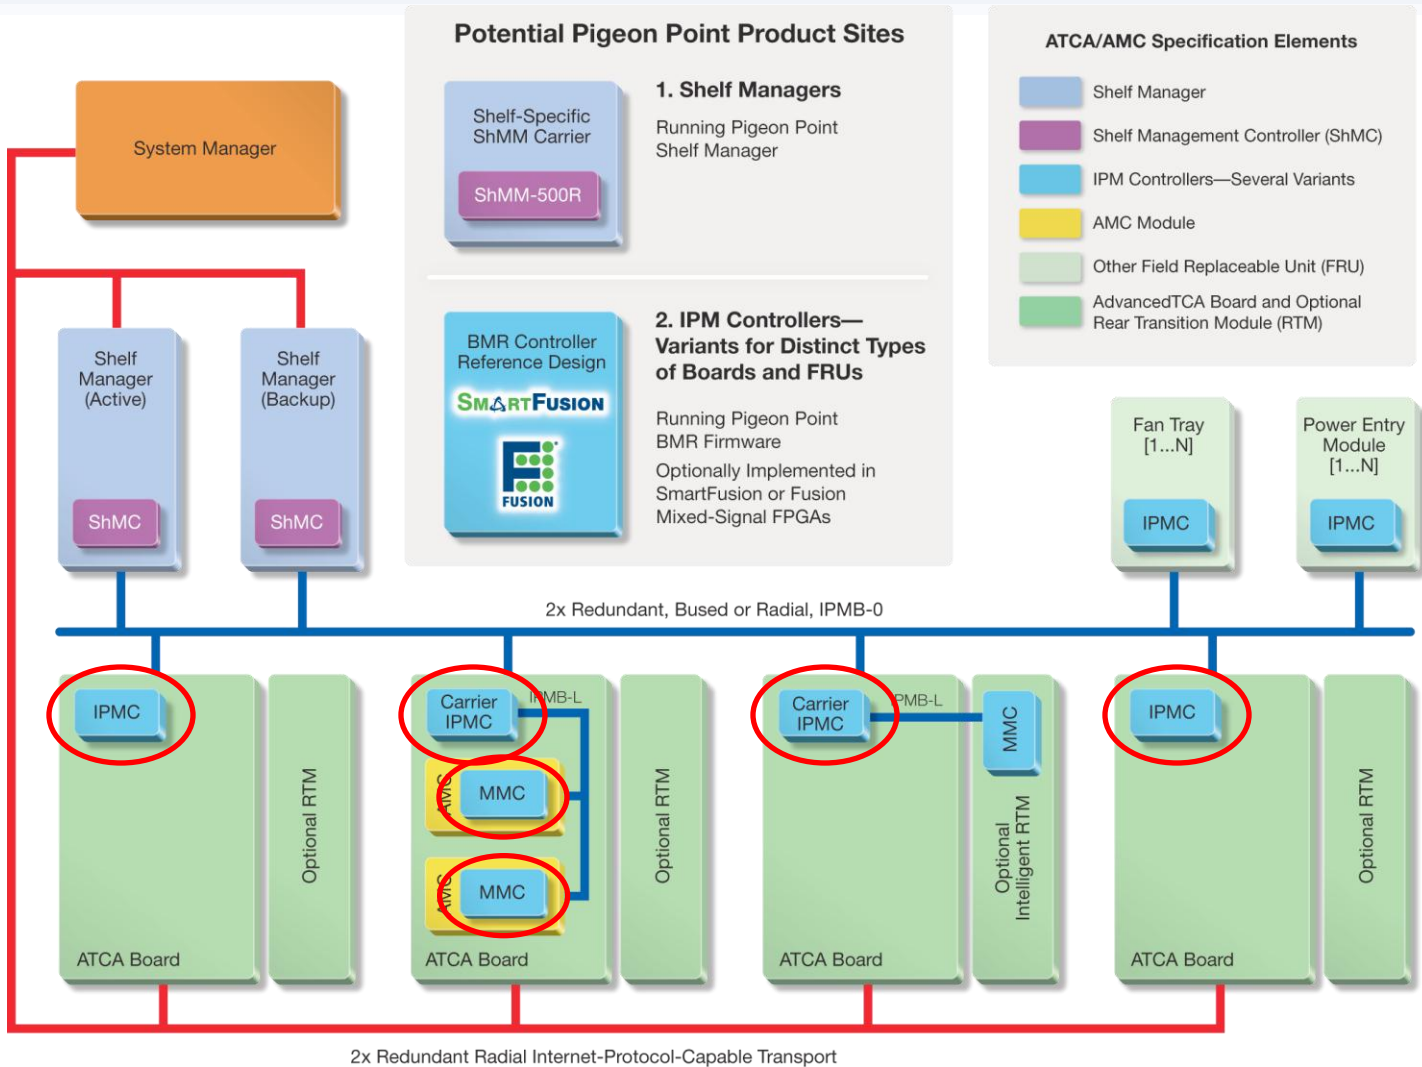
- Backup
 - Additional background on NC-SI architecture

Introduction to Pigeon Point Systems

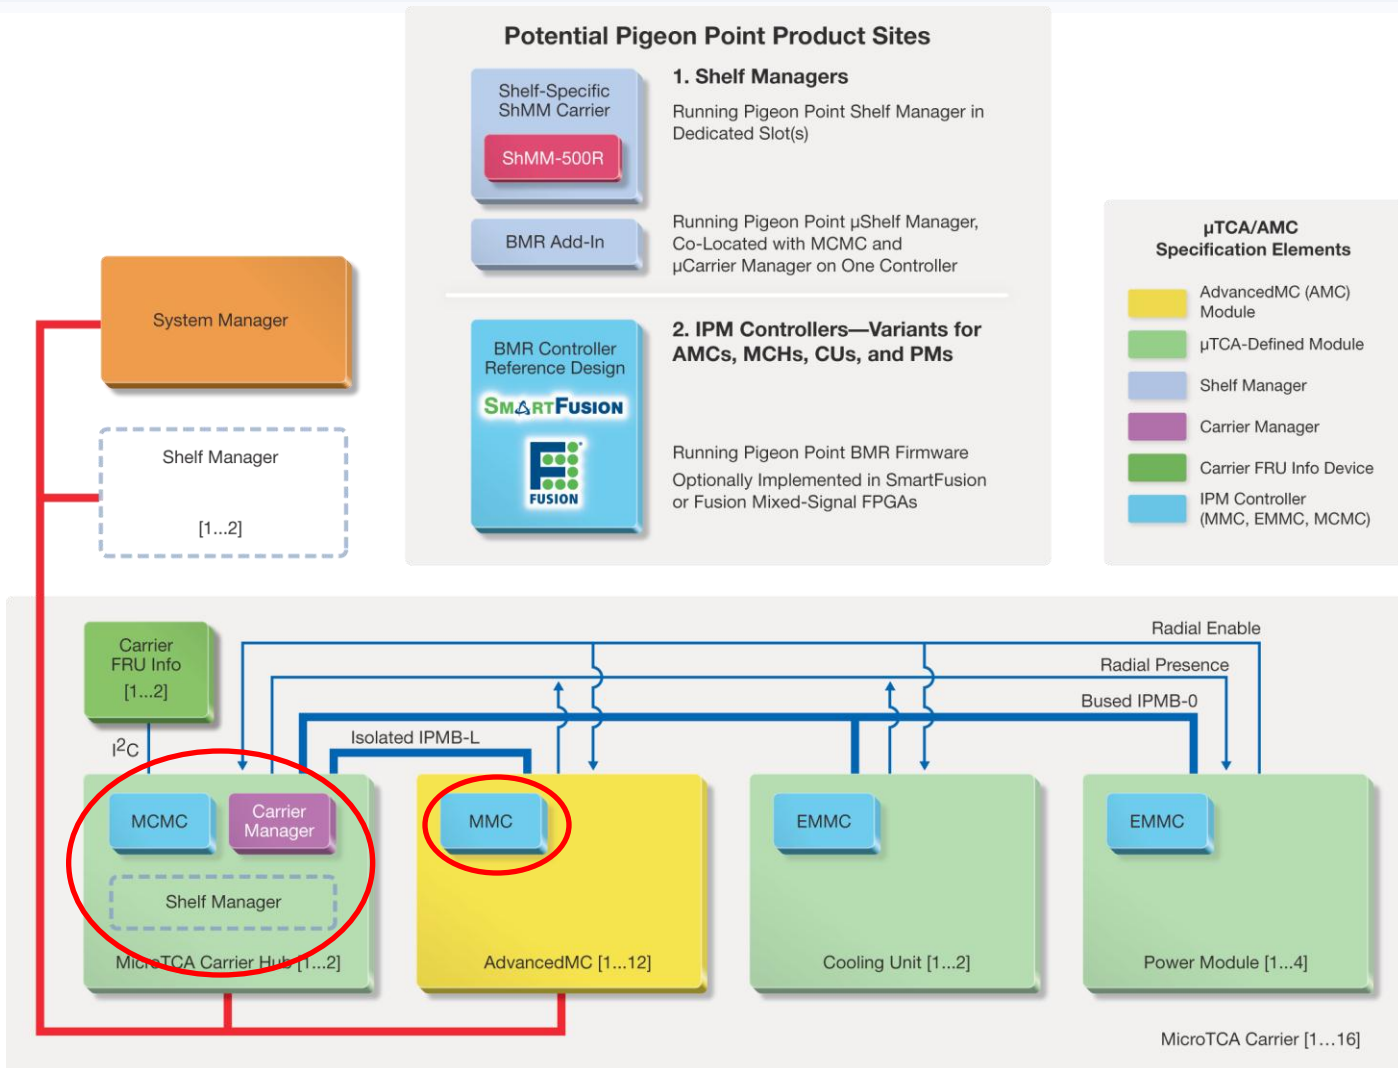


- Providing world-class management components for open modular platforms
 - Focused on delivering outstanding xTCA management controllers and Service Availability Forum interfaces to them
 - Dependable for specification compliance, interoperability and responsive support
 - Proven in over two hundred companies and with years of interoperability experience
- Founded in 1997 and headquartered in Scotts Valley, Calif.
- Acquired by Actel Corporation in July 2008
 - Actel was just acquired by Microsemi: Actel is now Microsemi
- Visit Booth 306 or www.pigeonpoint.com for more details

LAN-attached Management Controller (MC) Targets in ATCA/AMC

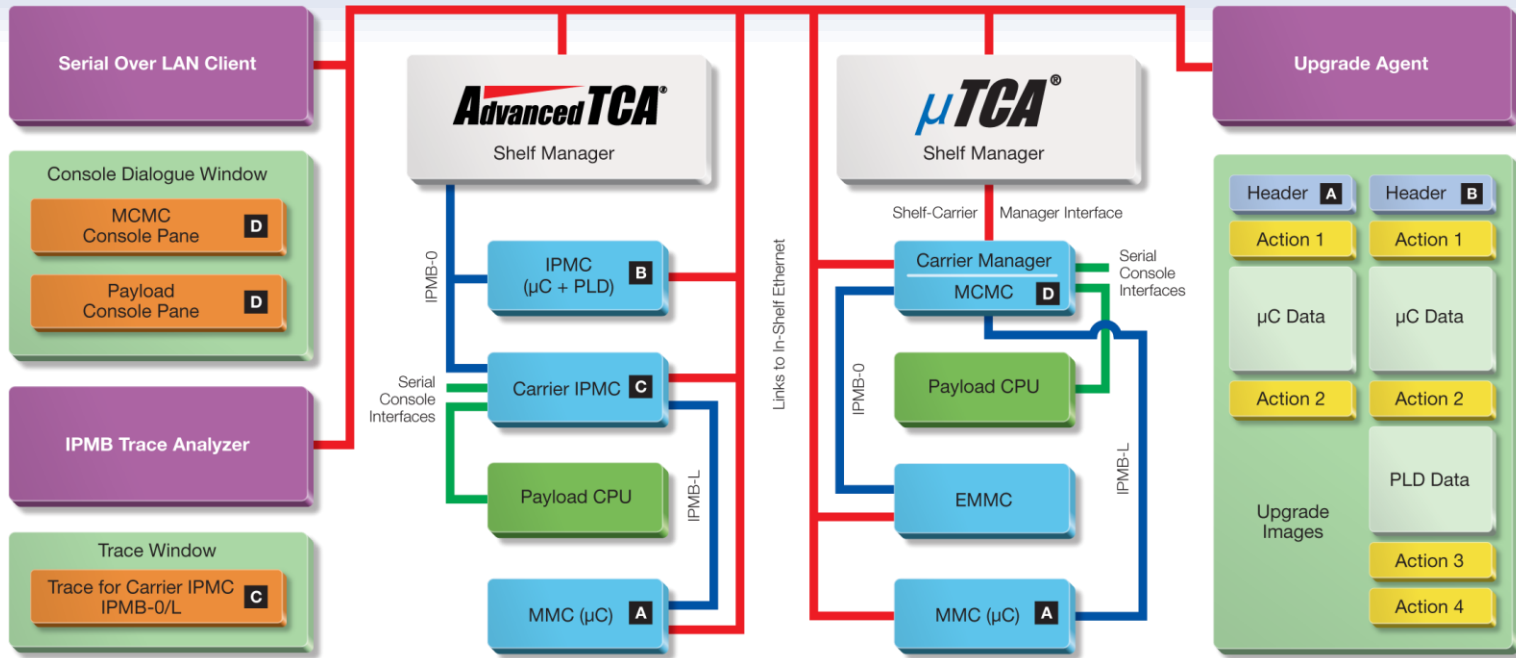


Main LAN-attached Management Controller (MC) Targets in μ TCA/AMC

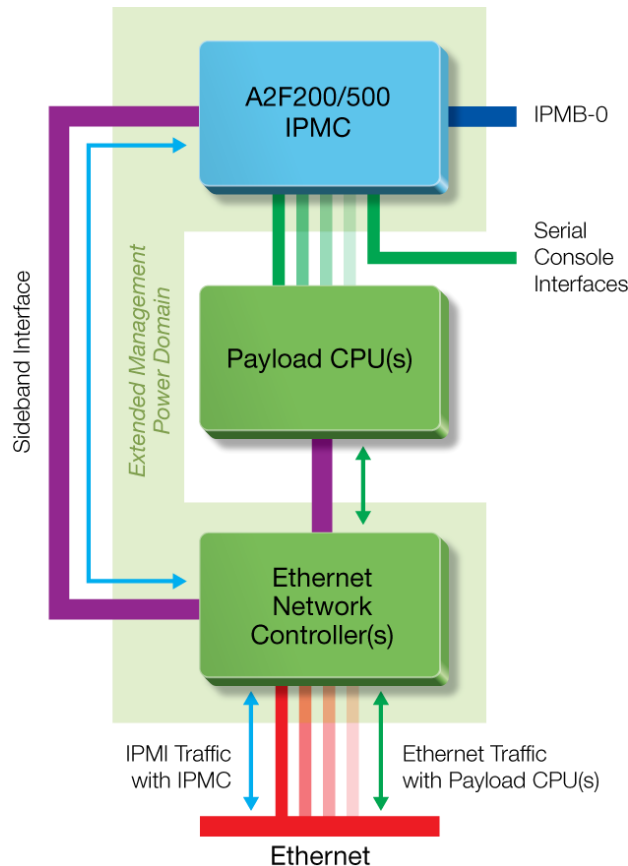


LAN-attached MCs: Motivations

Santa Clara, CA USA
 November 2010



- Serial over LAN (SOL) uses LAN for serial console traffic to avoid routing serial cables to boards
- HPM.1 upgrade transfers over LAN can be 15X quicker, especially important for large PLD image upgrades
- IPMB trace information over LAN complements SOL



- IPMC shares access to network controller that mainly serves payload
 - Network controller diverts incoming and merges outgoing management traffic
- Management traffic can include general IPMI traffic, plus SOL sessions, HPM.1 upgrades
- IPMC can attach to payload serial port(s) and use SOL for its own serial I/O
- Network controller may support multiple links and may be in management power domain
- Traditional sideband connection is I²C-based, but vendor- and/or product-specific
- DMTF (Distributed Management Task Force, dmtf.org) defines generic NC-SI sideband
 - Ethernet RMII (Reduced Media Independent Interface) as physical link

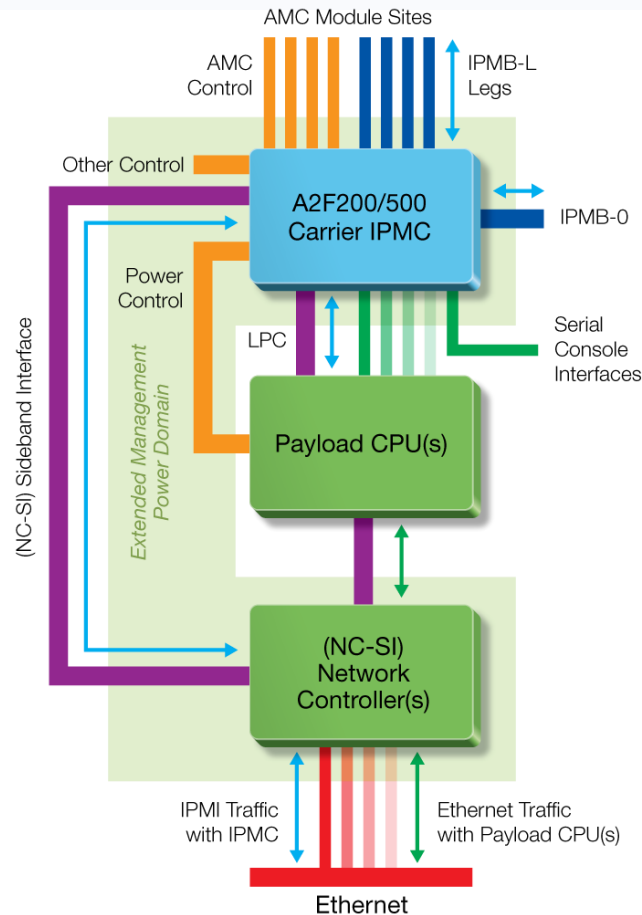
Key Benefits of MC LAN Attachment



- xTCA MCs are powered unconditionally and before payload
- If LAN access hardware is in the management power domain:
 - SOL accessibility for payload can include boot process from the beginning
 - MC can update PLD(s) that support payload before payload is otherwise powered/enabled
- SOL visibility for MC serial console can significantly aid diagnosing some tough problems in the field
 - Serial console cabling often not configured in the field
 - Some tough problems may require enabling special debug tracking output
- Remote visibility for IPMB interactions has similar benefits
- Higher bandwidth MC connection can yield big benefits in other areas, including firmware upgrades, access to large data areas

LAN Attach Can Be Especially Useful with Complex IPMC Architectures

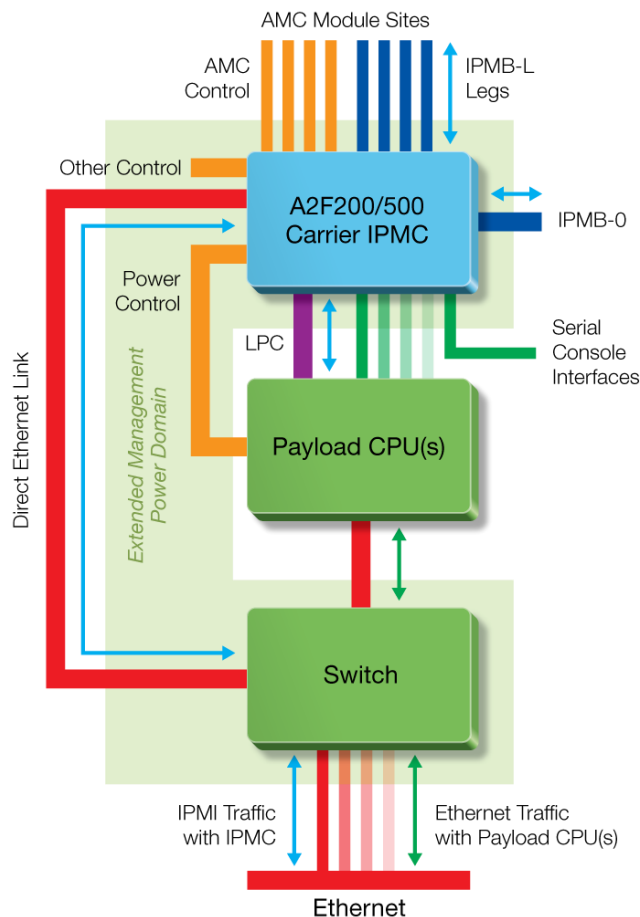
Santa Clara, CA USA
 November 2010



- Carrier IPMC can have multiple IPMB-L legs, increasing benefits of IPMB trace function
 - Understanding IPMB-L interactions can be important to diagnosing interop issues
- Complex board (say, a DSP farm) could have dozens or even hundreds of serial ports
 - Multi-drop serial could be used to simplify hardware aspects
 - Likely, only a modest number of ports need to be simultaneously visible via SOL
- Using IPMI over LAN channel for large MC data transfers can be big win over IPMB-0
- Intelligent mixed signal FPGA as IPMC device can integrate numerous other functions, including:
 - Low pin count (LPC) link to payload
 - AMC module and power control

LAN Attach Can Be Accomplished Via a Switch, as Well

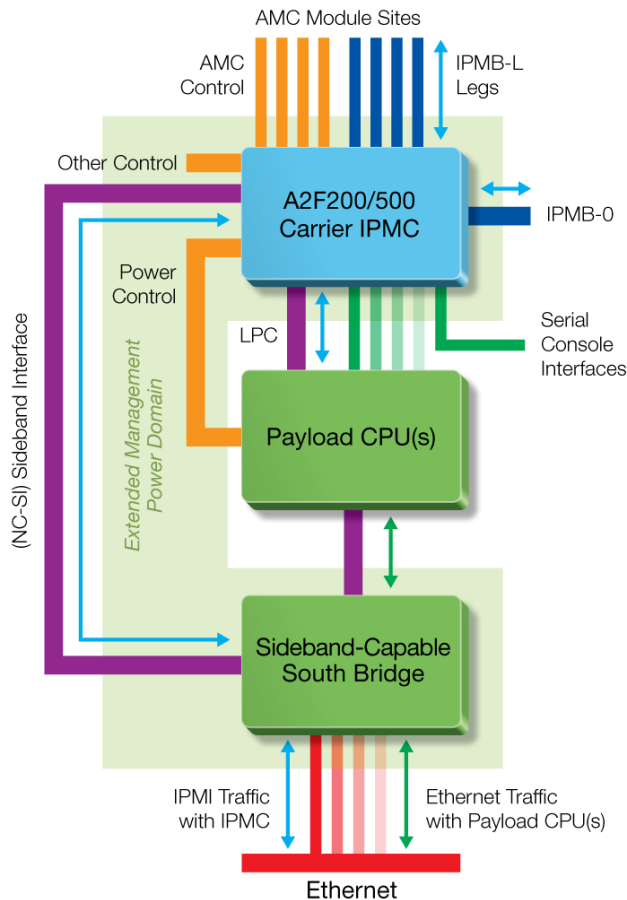
Santa Clara, CA USA
 November 2010



- Board could be an ATCA hub board, in which switch provides one or more fabrics for shelf
 - Payload would have on-board network connection to switch
- Alternatively, switch could provide Ethernet connectivity w/ AMC sites & in-shelf fabric(s)
- Independent of LAN attach device (switch or NC in these cases), extended management power domain may be a challenge
 - Fortunately, ATCA allows 15 Watts for board management power
 - Even so, detailed hardware issues with multi-domain power need to be solved
- Quite likely that in some cases management power will be insufficient
 - PICMG LAN attach specification aims to facilitate handling these cases


LAN Attach Can Also Be Realized via a Sideband-capable South Bridge

Santa Clara, CA USA
 November 2010



- Another variation, demonstrating the broad range of ways for realizing LAN attach
 - Included in reference design discussed by ADI Engineering in later session
- Sideband interface on south bridge may not be NC-SI, but proprietary I²C-based, instead
- Potential power challenges analogous to the LAN attach via switch case
- What if management controller is Module Management Controller (MMC) on an AMC?
 - Further challenge: management power for this case is spec-mandated ≤ 150 mA
 - Quite likely, south bridge or other LAN attach hardware must be payload powered in this case
 - Likely to benefit from spec assistance, such as a way to initially power just LAN attach hardware in payload


LAN Attach Facilitates Access to MCs, but Help w/ Parameter Assigns Needed



- Large systems can have dozens or hundreds of frames (racks) as complicated as example
- SOL, HPM.1 upgrades, IPMB trace, large FRU info fetches can be done via shared network connections
 - Could yield dramatic benefits for operational logistics
- But how do IP addresses and other low level parameters get assigned in scalable fashion?
- One answer: Dynamic Host Control Protocol (DHCP)
 - Could be configured to assign the same parameters to any board/module installed in a particular FRU position in system
 - Can further benefit operational logistics
- But, maximum benefit from using DHCP requires vendor-independent conventions on how DHCP clients request parameters

One Way to Implement LAN Attach Support in an ATCA (Carrier) IPMC

- Pigeon Point's Board Management Reference (BMR) solution for (Carrier) IPMCs using Actel SmartFusion mixed-signal FPGAs, including bench top hardware



PICMG Hardware Platform Management Extensions (HPM.x) Subcommittee



- Organized in June, 2010
- Includes numerous leading xTCA board and system developers
- Specific initiatives include:
 - Revising HPM.1 (IPMC Firmware Upgrade spec) to reflect experience since R1.0 (2007), coordination w/ specs below
 - Creating HPM.2 (LAN-attached IPMC spec), covering LAN attach arrangements for all xTCA board/module controllers
 - Creating HPM.3 (DHCP-assigned HPM Parameters spec) to simplify parameter assignments for LAN-attached IPMCs
 - Creating HPM.4 (Authenticating IPMC spec) to define a framework for IPMC-based “root of trust” on boards/modules
- Initial emphasis: HPM.2 and HPM.3

Key HPM.2 Goals in HPM.x Subcommittee

- Standardize arrangements for LAN attachments by xTCA board/module management controllers
 - LAN connections as optional supplements to mandatory IPMB connections
- Leverage appropriate foundation specifications, including relevant xTCA specs, plus IPMI 2.0 and NC-SI
- Include LAN-based support for:
 - Generic IPMI messaging, especially for commands that can involve large data transfers
 - Serial over LAN
 - HPM.1 upgrades
 - IPMI trace collection

Key HPM.2 Goals in HPM.x Subcommittee (2)

- Generic IPMI messaging
 - Support large IPMI messages to boost message throughput, especially for key commands can have large data payloads
 - Facilitate LAN attachment for challenging configurations, such as for [power w/ LAN interface] > [management power budget]
- Serial over LAN for boards/modules with up to ~250 physical serial ports that may need LAN based access
 - No plans to extend the IPMI 2.0 limit of 15 ports concurrently accessible via Serial over LAN
- HPM.1 firmware upgrades (mainly ensuring that HPM.1 upgrades can take advantage of large message sizes)
- IPMI trace access via LAN-resident client that collects trace packets and forwards to appropriate analyzer application

Key HPM.3 Goals in HPM.x Subcommittee (2)



- Define provisions for DHCP assignment of key hardware platform management (HPM) parameters
 - Scalable and standardized parameter assignment methods complementing implementation-specific approaches
 - Example: standardized way to identify any FRU position in an xTCA shelf, to facilitate DHCP parameter request
- Target scope
 - Network-related parameters associated with xTCA Shelf Managers
 - Network-related parameters associated with xTCA LAN-attached board and module controllers
 - Specific non-network-related parameters, plus any other parameter defined for IPMI LAN configuration command

Conclusion: LAN-attached xTCA Management Controllers (MCs)...



- Can significantly simplify operational logistics and reduce cost for serial console access in large xTCA systems
- Can dramatically improve performance for key MC-related operations, including firmware updates, large FRU info fetches
- Can share the use of a payload-focused in-shelf LAN
- Benefit from standardization of the sideband interface via the NC-SI specification from DMTF
- Are gaining increasingly sophisticated support, including in widely used Pigeon Point BMR management controller solutions
- Are especially facilitated by Actel SmartFusion intelligent mixed-signal FPGA based MCs that:
 - Include RMI 10/100 Ethernet MAC, natively supporting NC-SI attachments to NCs
 - Can run 32-bit ARM processor at 80 MHz or more, to provide needed processing bandwidth

Thank you!!

Mark Overgaard

mark@pigeonpoint.com



World-Class Management Components

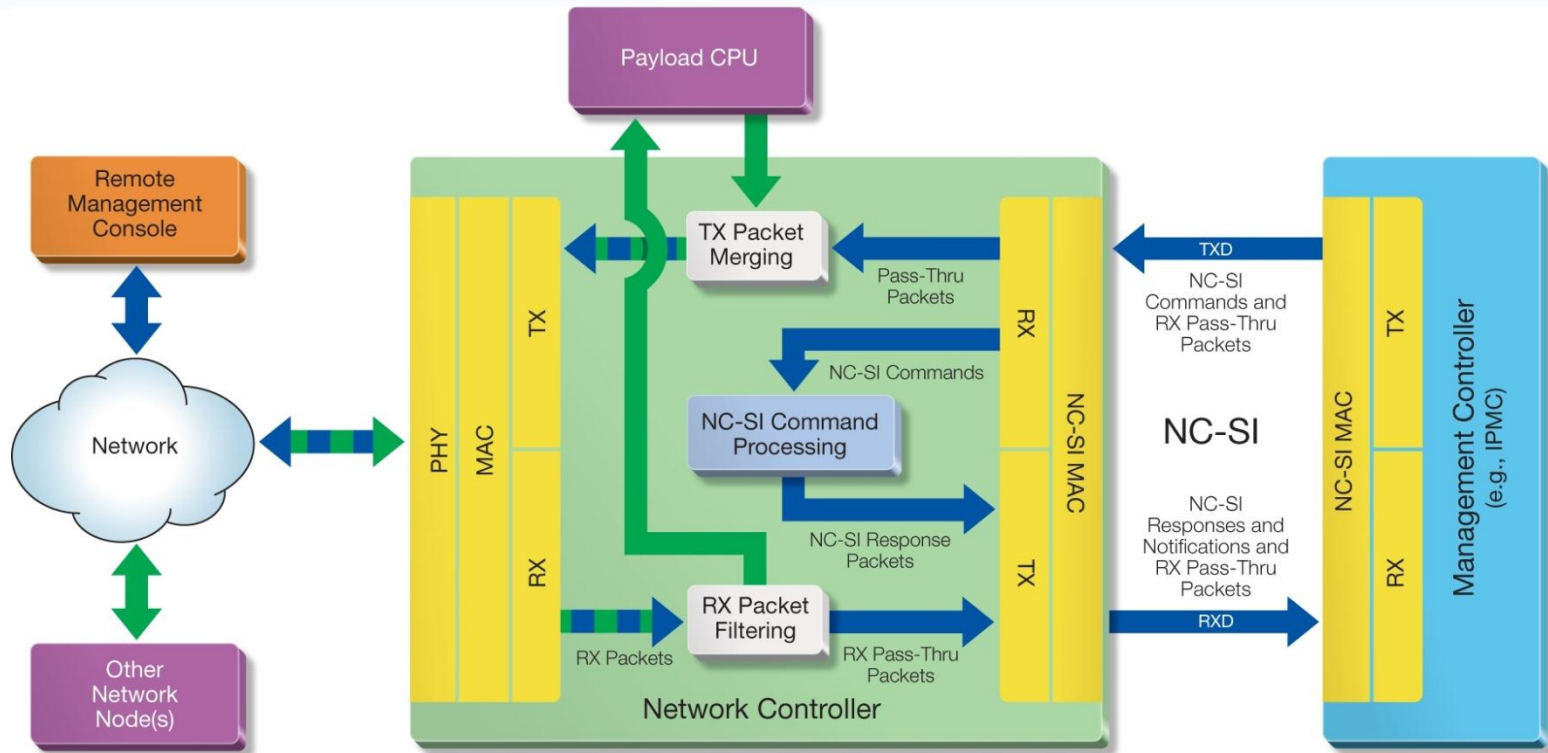
FOCUSED. DEPENDABLE. PROVEN.



Mark Overgaard is the founder and president of Pigeon Point Systems (PPS), which is now an Actel company. Mr. Overgaard is also a leader in the technical subcommittees of PICMG, including the hardware platform management aspects of AdvancedTCA, AdvancedMC and MicroTCA. Prior to founding PPS in 1997, Mr. Overgaard was VP, Engineering at Lynx Real-Time Systems (a Unix-compatible RTOS supplier) and TeleSoft (a major supplier of embedded development solutions for Ada).

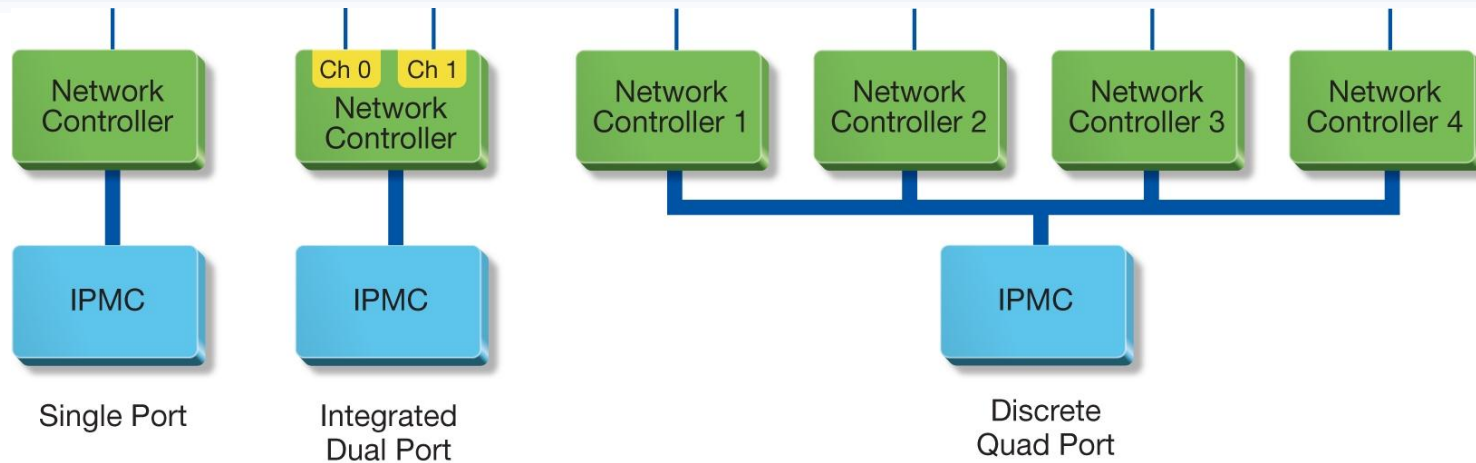
- Further background on NC-SI, including overviews of:
 - Internal operations within an NC-SI-capable NC
 - LAN attach options facilitated by NC-SI specification

How Does NC-SI Work?



- Filtering and merging inside NC allows shared use of external LAN
- NC-SI specification at www.dmtf.org has full detail, including commands and protocols for MC-NC interactions

Some Advanced NC-SI Configurations



- NC-SI architecture provides for advanced configurations beyond simple single port
- Integrated multi-port NCs (such as widely available dual port GbE or 10GbE variants or even new four port NCs)
- Multi-drop configurations with up to four distinct NCs on a single RMI link; arbitration alternatives:
 - Hardware arbitration via serial opcodes around a ring of NCs
 - Command-based arbitration controlled by IPMC