

Standardized LAN-Attached Management Controllers Yield xTCA Performance and Serviceability Gains

LAN-attached management is leading to faster and richer system management applications in ATCA and MicroTCA systems, both of which also use ATCA cards.

by Mark Overgaard, Pigeon Point Systems

Every board and module built for the popular xTCA open modular ecosystem must have a local management controller, and each controller must have an I²C-based link to its management parent. These management controllers and the I²C “network” that links them are crucial to the widespread adoption of xTCA in telecommunications and (increasingly) in other applications, which already drive a half billion dollar annual xTCA market, as AdvancedTCA (ATCA), the dominant segment of xTCA, enters its second decade. xTCA also includes MicroTCA (μ TCA) and AdvancedMC (AMC), a hot swappable module architecture used in both ATCA and μ TCA systems.

At the same time, most of those boards and modules—collectively, field replaceable units or FRUs—also con-

nect to a much faster in-shelf serial fabric, often Ethernet. LAN-attached xTCA controllers connect to such a fabric to supplement the much slower I²C link. “LAN-Attached TCA Management Controllers: How to Build and Use Them” in the October 2009 issue of *RTC*, introduces this concept, which is already widely used in xTCA products. The news now is that LAN attachment is standardized for xTCA controllers, with the recent adoption of the PICMG Hardware Platform Management (HPM) HPM.2 and HPM.3 specifications. Pigeon Point Systems proposed and chairs the PICMG subcommittee that has developed these specifications.

Figure 1 shows the LAN-attached management framework for ATCA and μ TCA, along with ways to take advantage of that framework. Without the management controller LAN connections, external network clients are typically limited to interacting with management controllers indirectly via a shelf manager.

What Applications Benefit?

Two of the applications shown in Figure 1 were introduced in the October 2009 article referenced above: firmware upgrades and serial over LAN. What new facilities in these areas result from the adoption of the Hardware Platform Management (HPM) HPM.2 and HPM.3 specifications? Both benefit from the overall strengthening of the LAN attach infrastructure in HPM.2 and the standardized IP address management facilities of HPM.3.

For firmware upgrades, the most important new development is a standardized way to support the large Intelligent Platform Management Interface (IPMI) messages (hundreds of bytes in size) that are possible over Ethernet, compared to the 32 byte limit for Intelligent Platform Management Bus (IPMB) messages. This difference can contribute 20 percent or more to the 10x or more faster upgrades that are possible with a LAN attach connection.



Get Connected

with companies mentioned in this article.

www.rtc magazine.com/getconnected

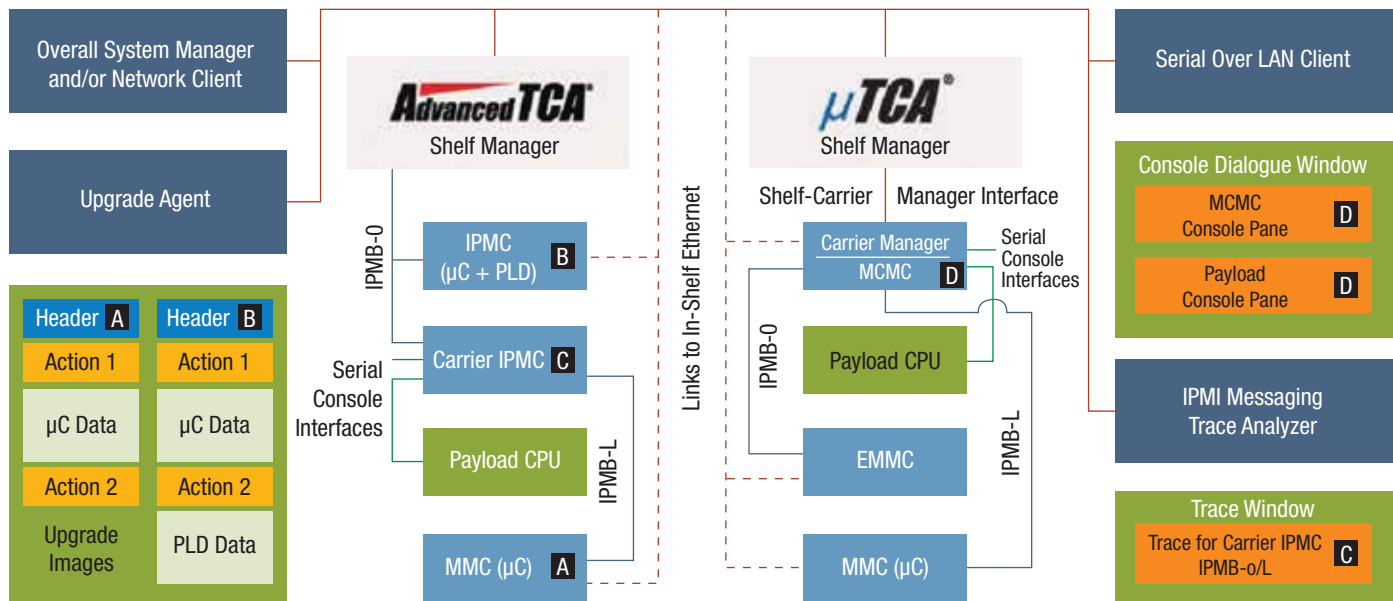


FIGURE 1

LAN-attached management framework for ATCA and μ TCA shows applications for management controller LAN connections (the dashed line links).

The entire serial over LAN facility is more robust in HPM.2, including support for over 250 serial ports per FRU and a mechanism for a network client to get a list of the available ports on the board.

The detailed architectural provisions in HPM.2 facilitate additional applications, such as direct interactions with a management controller by an overall system manager or other network client. An xTCA system manager can monitor dozens or even hundreds of shelves. Most of the system manager's interactions occur with the shelf managers of those shelves to get shelf level summary information. Some system managers also rely on frequent interactions with lower level controllers. For instance, when a new shelf goes under management, the system manager may need to fetch thousands of bytes of descriptive information about its controllers; doing that over LAN can make a dramatic difference.

Then there are shelf manager interactions. Most of the messages between a shelf manager and one of its controllers are engineered for the low transfer rate of IPMB, but sometimes big or especially numerous transfers are necessary. For instance, when a new shelf starts up, the shelf manager may need to retrieve descriptions for thousands of sensors and field replaceable units (FRUs) in the shelf. Doing such retrieval via LAN can drastically reduce the start-up time, making the shelf usable and fully manageable earlier. Furthermore, even though the routine messaging in a steady state shelf can be handled fine by IPMB, important messages can be seriously slowed during a period of high IPMB usage for other purposes, such as firmware upgrades or heavy system manager interactions. Both of these applications are dependent on HPM.2 compliance in the lower-level controllers and are not really feasible without such standardization.

Another class of LAN attach applications is simply not practical with only IPMB connections to a shelf's controllers. This class includes the serial over LAN application mentioned above, in addition to IPMI message tracing. xTCA management is based on the IPMI messaging architecture and those messages can flow across a LAN connection as well as an IPMB. In either case, for system diagnosis, having access to traces of such messages can be highly beneficial; consider, for instance, the IPMB-L traffic among a Carrier IPMC like controller C in Figure 1 and its subsidiary (type A) MMC(s). Clearly, attempting to provide a trace of IPMB traffic across IPMB itself would be folly, so a LAN attach connection is vital for IPMB traffic diagnosis via traces.

xTCA controllers and these applications implemented in accordance with HPM.2 and HPM.3 can be from different vendors and even completely different implementations and still be interoperable.

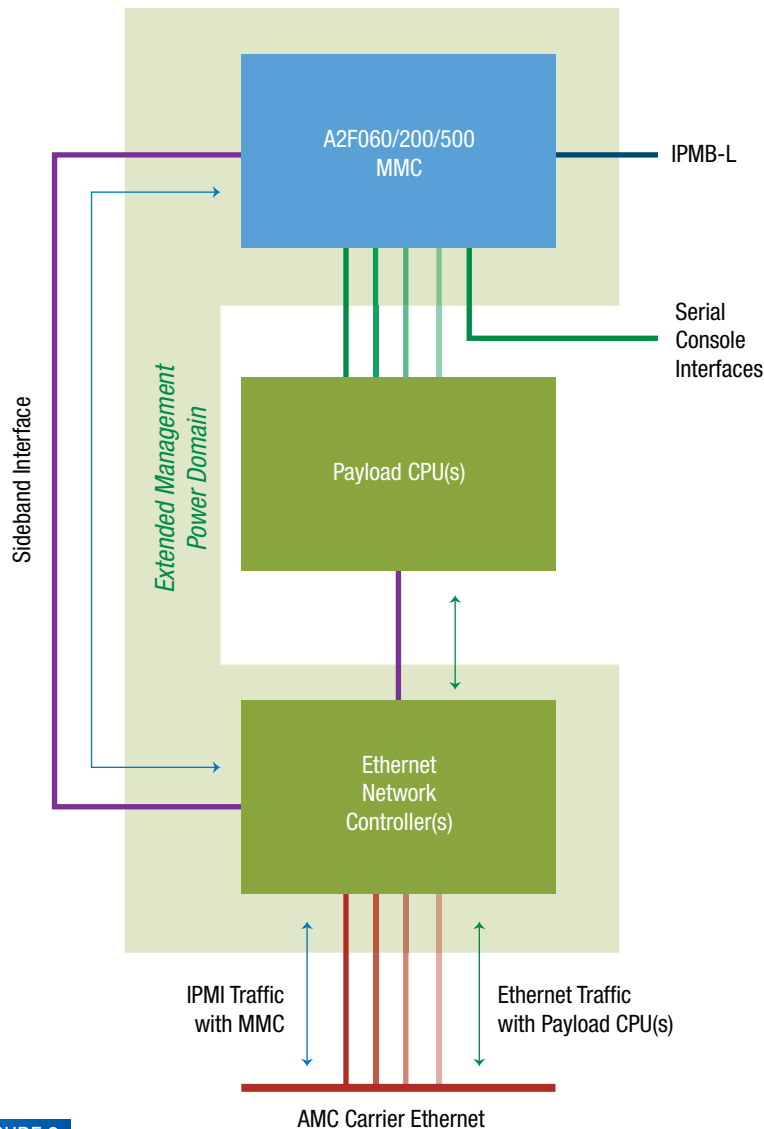


FIGURE 2

LAN-attached MMC connected via sideband interface to shared network controller.

What Is HPM.2 and What Does it Cover?

PICMG HPM.2, the LAN-Attached IPM Controller specification, covers LAN attachment for any xTCA board or module controller and uses “IPM Controller” as a generic term to stand for all the xTCA FRU controller types.

Figure 2 shows a simple LAN-attached controller, in this case an MMC. In this and the Figure 3 example, the controller is based on Microsemi Corporation’s SmartFusion intelligent mixed signal FPGA. This FPGA comes in several models, A2F060 through A2F500, which are primarily distinguished by their logic capacity. “Customizing a Microcontroller

for Hardware Platform Management” in the June 2011 issue of *RTC* provides more background on SmartFusion-based management controllers.

In the MMC example, there are Ethernet controller(s) already on the AMC to provide connectivity for the payload CPU(s) to an Ethernet fabric implemented on the AMC carrier. The MMC shares access to that Ethernet connectivity via a sideband interface implemented in the network controller, which combines IPMI traffic with the MMC and generic Ethernet traffic with the payload CPUs so that no dedicated management Ethernet is needed to support direct interactions with the MMC. The June 2011 *RTC* article ref-

erenced above provides more details on sideband interfaces, including a widely implemented open standard sideband interface called NC-SI that effectively offers a 10/100 Mbit/s Ethernet connection to the network controller for the management controller.

HPM.2 also explicitly covers the topic of extending the management power domain, as shown in Figure 2, to include the network controller in addition to the MMC, so that interactions with the MMC are fully viable, even if the payload CPU complex is unpowered or faulted. The AMC specification limits 3.3V management power to 150 mA, which is likely not enough for both the MMC and a gigabit class network controller. Therefore, an implementation like this probably has to supplement the management power with a portion of the 12V payload power supplied to the AMC by its carrier.

Another similar challenge has to do with the electronic keying (e-keying) that xTCA uses to manage carrier and backplane links, such as the AMC carrier Ethernet(s) in this example. Normally xTCA e-keyed links are not enabled unless the payload is operational, but that would significantly reduce the utility of links used for management, since many of the most important diagnostic and performance benefits happen with the payload off either purposely or due to a fault.

Figure 3 shows a more complex LAN-attached controller example: a Carrier IPMC (IPM Controller) that implements local management for an ATCA board that provides AMC slots (four of them in the example). Again, the management Ethernet traffic is shared with payload-oriented traffic, but also in this case with AMC module traffic, all enabled by an Ethernet switch that connects with the ATCA backplane’s Base Interface Ethernet. Conveniently, a management controller that supports an Ethernet-based sideband interface can use exactly the same port to make a direct connection with an Ethernet switch, as in this example.

Here, the extended management power domain needs to include the base interface switch, which may require substantial power. Fortunately, the management power budget for an ATCA IPM

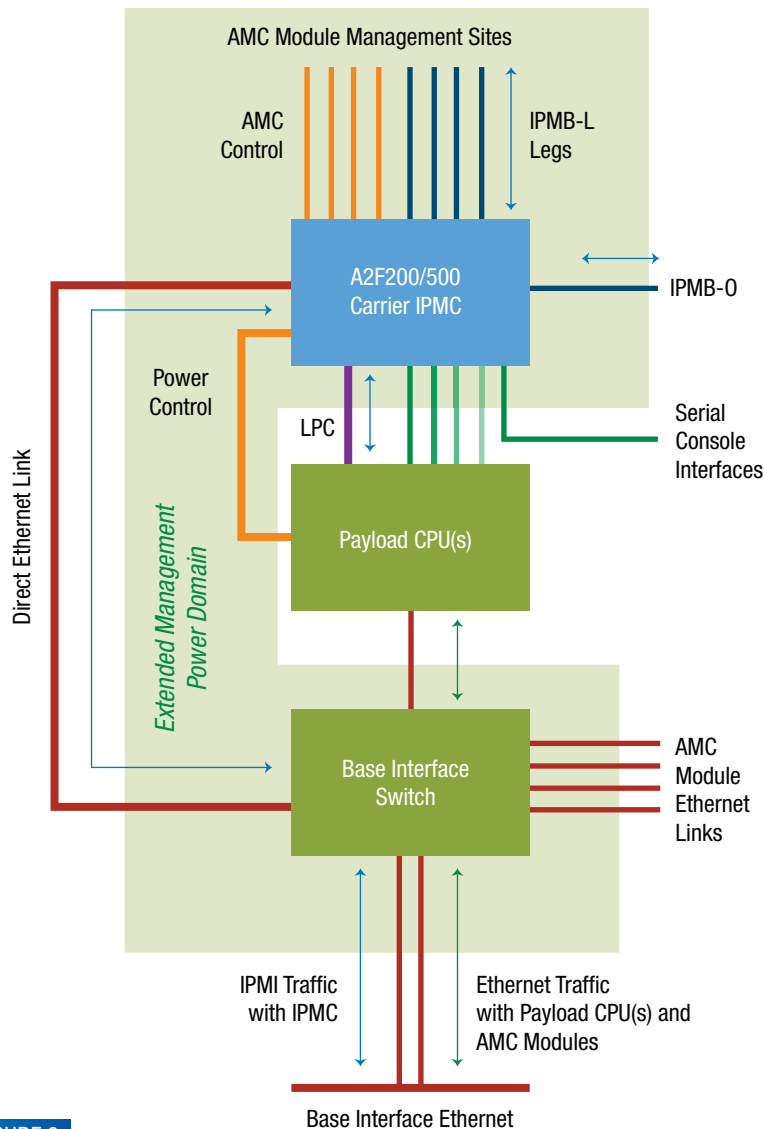


FIGURE 3

LAN-attached Carrier IPMC connected via shared Ethernet switch.

Controller is 15W, which may be enough. If not, some of the payload power budget must be used to supplement management power, and power budget negotiations with the shelf manager for this board must reflect that.

Furthermore, the extended management power domain in this example extends even to the MMCs on AMC modules hosted by this carrier, as shown in the figure. The likely need for 12V payload power for LAN attach connections on those modules adds to the AMC carrier power implementation and budgeting challenge.

Supporting the above facilities in an implementation-independent and interoperable fashion takes a considerable por-

tion of the HPM.2 specification, which also covers additional advanced features, such as ways for a network client to take advantage of redundancy in network paths to a management controller.

Where Does HPM.3 Fit In?

PICMG HPM.3, the DHCP-Assigned Platform Management Parameters specification, addresses an additional challenge not mentioned so far: how do all those management controllers get Internet Protocol (IP) addresses assigned so they can communicate over the in-shelf LANs? Every IP address has to be unique within a given network domain, and a domain (such as a large mobile switching center) could have thousands of management



Embedding Excellence

MSC Embedded offers a large variety of COM products in COM Express™, Qseven™ and ETX®. Our COM products are scalable, easy to install and powerful to enable long living industrial applications with an upgrade path to future needs.

COM Express™ MSC CXB-6S


Intel® Core™ – 2nd Generation

- Multiple i5, i7 and Celeron processor options available
- Single, dual and quad core solutions
- Intel® HD Graphics 2000 / 3000
- Intel® QM67 or HM65 Platform Controller Hub
- Up to 16GB DDR3 SDRAM, dual channel
- Four SATA-300 interfaces
- One PATA interface
- LVDS (24 Bit, dual channel) and VGA
- Resolution up to 2560 x 1600
- Five PCI Express™ x1 lanes
- Eight USB 2.0 interfaces
- COM Express Type 2
125 mm x 95 mm [4.92 x 3.74"]

V-2_2012-WOEI-5945_RTC



MSC Embedded Inc.
direct: (650) 616 4068
info@mscembedded.com
www.mscembedded.com



OceanServer Digital Compass Products:

- Low Cost Family of Electronic Compasses
- Solid State Package
- High Accuracy Offered in Serial, USB or TTL
- Under \$200.00 in Low Volume
- Hard & Soft Iron Calibration
- Fully Engineered Solution For Embedded Applications

OceanServer
TECHNOLOGY, INC.

ocean-server.com
(508)-678-0550

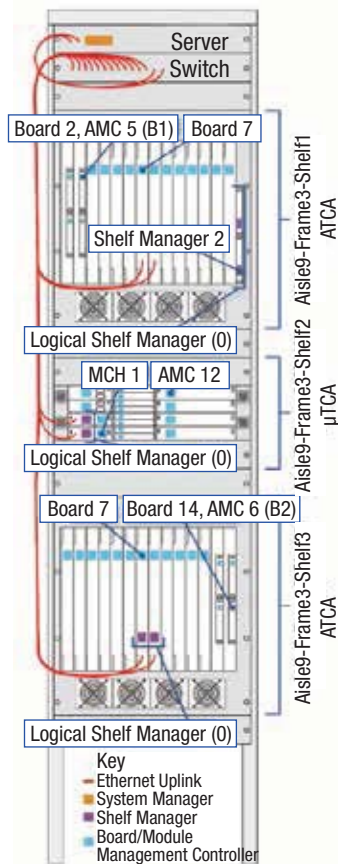


FIGURE 4

Example xTCA frame showing representative HPM.3-style FRU identifiers.

controllers. Figure 4 shows an example frame or rack.

One key to the puzzle is the Dynamic Host Configuration Protocol (DHCP), an Internet standard for managing IP addresses and related networking parameters. In DHCP, one or more servers in a network domain are in charge of dispensing network parameters, including IP addresses; clients in the domain use a standardized protocol to discover these servers and request parameters. In order to ensure that each such client, such as a LAN-attached management controller, gets a unique set of parameters, each client must provide a unique identifier.

HPM.3 defines a vendor- and implementation-independent way to identify client management controllers that allows uniqueness to be ensured, while also allowing the data base to be set up just once for a given installation, with no changes

required when boards, modules or even entire shelves are replaced. In HPM.3, management controllers identify themselves by their location.

The first part of the location is a shelf address, which could be geographic coordinates that designate the aisle, frame and shelf position. When a frame of xTCA shelves, like the example in Figure 4, is provisioned or installed, non-volatile storage in each of the shelves is updated with the coordinates that uniquely identify it. If a given shelf is replaced, the replacement shelf gets the same coordinates and uniqueness is preserved. Each of the shelves in Figure 4 is labeled with a simplified shelf address.

The second part of the location is the coordinates of the board or module within a shelf. xTCA FRUs can determine those coordinates from the slot in which they're installed, so a board or module replacement doesn't require any reconfiguration at all. The replacement board identifies itself indistinguishably from its predecessor in the same slot and gets the same address assignment from DHCP. Figure 4 shows some simplified examples of FRU identifiers. In the example, several of the ATCA boards are AMC carriers; the two AMC slots that each of them implement, according to the AMC specification, are designated B1 and B2, with corresponding numeric values of 5 and 6. Also, there are several instances of logical shelf manager identifiers; these are used to get IP address assignments for the main interface used by a system manager to connect with a shelf, independent of which of the two redundant shelf manager instances is active.

Location-based identification for xTCA FRUs fits well with a frequently used approach to xTCA system architecture, where specific application responsibilities are assigned to particular slots or sets of slots in shelves and the boards in those slots handle those responsibilities. Assigning IP addresses in the same way can be very convenient. ▀

Pigeon Point Systems
Oceanside, CA.
(760) 757-2304.
[www.pigeonpoint.com].