



AdvancedTCA specification revision adds hardware platform management benefits

On May 19 2006, PICMG adopted a substantial revision of the AdvancedTCA specification, reflecting more than three years and thousands of AdvancedTCA shelves worth of experience. In this column Mark highlights this revision in the hardware platform management arena. This second revision affecting R2.0 of the AdvancedTCA specification is known informally as ECN-002.

Hardware platform management focuses on fundamental aspects of the AdvancedTCA hardware, including collecting inventory data, budgeting power, tracking and responding to temperature events, and brokering high-speed fabric connections. In the initial AdvancedTCA specification, this area was known as *shelf management*, but that term is now used primarily for the shelf-wide level, which consolidates information and control for all the Field Replaceable Units (FRUs) or boards in a shelf.

Key improvements ECN-002 brings to the table include more effective and portable system managers, finer grain thermal management, and more general topologies for radial IPMB-0.

More effective and portable system managers

One theme of the extensions is providing more visibility to the system manager for decisions made by the shelf manager, such as in the power and cooling areas. Some system managers display the state of the shelves in a network operations center with graphics; with the ECN facilities, these displays can include more detailed information.

Another theme is improving robustness of shelf/system manager interactions. One extension ensures that multiple independent components in the system manager don't attempt to concurrently modify the shelf configuration information, possibly wreaking havoc in a shelf. Another extension informs the system manager of all the Internet Protocol (IP) addresses with which the shelf manager and its redundant

instances can be accessed. This allows the system manager to continually monitor the health of all these interfaces and initiate recovery actions if any of them go bad. On the cooling front, when a fan-related problem is being diagnosed, the system manager can temporarily remove that fan from the shelf manager's control. The system manager can then exercise that fan thoroughly without conflicting with the shelf manager.

Finer grain thermal management

Prior to ECN-002, shelf managers could control the fan speeds of different fans implemented in the shelf via generic PICMG-defined IPMI commands. However, without ECN-002, shelf managers could not discover in a standardized way which fans would affect specific FRUs' cooling states. Therefore, if any FRU reported a high-temperature exception, all the fan speeds had to be increased.

Deployed shelves must minimize the extra noise that higher fan speeds produce. With a shelf architected to take advantage of ECN-002's fan geography facility, the response to a high-temperature exception from an FRU can be an increase in the speeds of the fans for just the zone where that FRU is installed, not all the fans. The 14-slot shelf in Figure 1, for example, has three cooling zones. Slots 1-5 are in cooling Zone 1, with slots 6-10 and 11-14 in Zones 2 and 3 respectively. The pair of dedicated shelf manager slots on the right side of Figure 1 is also part of Zone 3.

A high-temperature exception in Slot 7 would result in boosted fan speeds for only Zone 2 (one third of the fans in the shelf, not all of them). Such a shelf could even have a separate fourth zone to add even more control granularity for the Rear Transition Modules (RTMs) that can be installed behind the front boards.

A key aspect of this facility is that the shelf describes its own fan geography as part of its shelf configuration information. Therefore, if the shelf managers in the example shelf were extracted and moved




Figure 1

to a different shelf with completely different fan geography, they could automatically adapt to the new configuration.

More general topologies for radial IPMB-0

IPMB-0 is the main communication mechanism inside the shelf for hardware platform management messages. Some AdvancedTCA users consider a radial implementation of IPMB-0 to be mandatory for certain types of deployments. As shown in Figure 2, a radial implementation can have separate IPMB-0 segments going to the IPM controller on the board in each slot. This keeps problems on one slot from affecting the other slots. Consider, for instance, the possibility of a board malfunction that brings down IPMB-0 in a remote shelf. If IPMB-0 is implemented on a radial basis, overall operation of the shelf can be preserved by disabling the IPMB-0 segment for that slot until service personnel can get to the shelf and replace the board.

On the other hand, proponents of implementing IPMB-0 on a simple based or multidrop basis note that IPMB-0 is dual-redundant, with IPMB-A and IPMB-B to back each other up. If a blocking problem

happens on IPMB-A, IPMB-B can be used until repair can be accomplished. These designers believe that the extra complexity and cost of a radial IPMB-0 implementation is not cost-effective.

Even among radial IPMB-0 enthusiasts, however, there are different favorite topologies. In one widely used topology, as shown on the left side of Figure 2, each dedicated shelf manager has a radial hub for both IPMB-A and IPMB-B. The distinct radial segments from each of these hubs come together at the A and B ports of the IPM controllers. This topology preserves dual segment access to every IPM controller even when only one shelf manager is present or operational.

Another widely used topology, as shown on the right side of Figure 2, implements an IPMB-A hub on one shelf manager and an IPMB-B hub on the other shelf manager. In this topology, both shelf managers need to be present and operational to have dual-segment access to all IPM controllers.

Prior to ECN-002, the shelf configuration records could not describe the right-hand topology in Figure 2 because the data structures in the AdvancedTCA specification had been conceived primarily to describe variations of the left-hand topology. With ECN-002, both these topologies and a wide range of others can be described, so that a portable shelf manager can automatically adapt to the specifics of a given shelf.

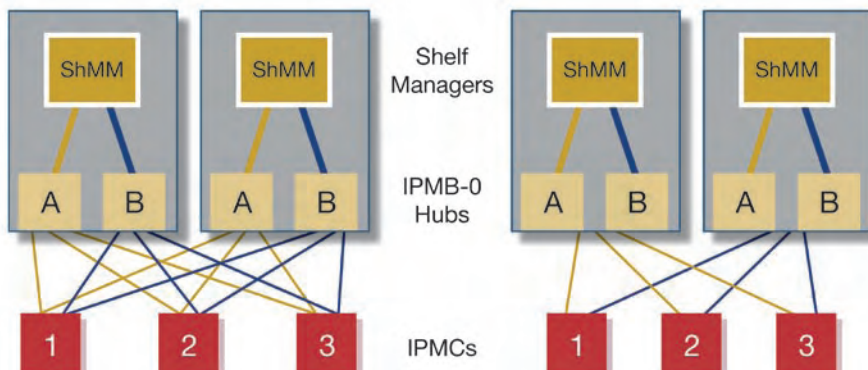


Figure 2

Key New Commands	New System Manager Capabilities	Resulting System Manager Benefits
Set/Get Fan Policy	Discover the fan to FRU mappings Discover and control which fan devices are enabled for autonomous Shelf Manager control	Getting fan to FRU mappings without having to parse detailed records in shelf information Ability to temporarily disable shelf manager autonomous control of particular fan devices, say for diagnostic operations
Get Shelf Power Allocation	Discover the power allocations that have been made by the Shelf Manager for all power feeds	Visibility on allocation decisions made by the Shelf Manager
Get Shelf Manager IP Addresses	Discover all the IP addresses by which the Shelf Manager and its redundant instances can be contacted	This collection of addresses can be routinely monitored, with automatic follow-up if any address stops responding
FRU Inventory Device Lock Control and Write	Acquire and use exclusive access to shelf configuration information	Ensuring that concurrent attempts to change shelf configuration information do not cause corruption thereof

Table 1

Shelf manager delivers first support for ECN-002

One example of a shelf manager that complies with ECN-002 is the widely used Pigeon Point Shelf Manager Release 2.3.0, which is shipped on an SODIMM-sized shelf management mezzanine, such as the second generation ShMM-500R shown in Figure 3.



Figure 3

A generic shelf manager delivered by Pigeon Point in the flash memory of an ShMM-500R can automatically adapt

across a wide range of shelf architectures by looking at the shelf configuration data:

- Bused IPMB-0 or radial IPMB-0 and within radial, a wide range of topologies
- Fan geographies (that is, fan-to-FRU mappings) that are completely undescribed, as in pre-ECN-002 shelves, to a wide range of fine-grained descriptions enabled by the ECN
- The number of slots and the high-speed fabric topologies implemented in those slots, which were already self-describing prior to ECN-002

The additional facilities in an ECN-002-compliant system manager interface support the needs of real system manager applications and can be used with any ECN-002-compliant shelf manager from any source to protect the Telecom Equipment Manu-

facturers' and their customers' investment in the often complex system manager layer. Table 1 summarizes key system manager related extensions in ECN-002.

Mark Overgaard founded Pigeon Point Systems, a leading supplier of AdvancedTCA, AdvancedMC, and MicroTCA management components, in 1997 to focus on products and services supporting the adoption of open modular platforms to replace proprietary architectures. Mark earned an MS in Computer Science from UC San Diego and a BS in Physics from Geneva College.

To learn more, contact Mark at:

Pigeon Point Systems

P.O. Box 66989
Scotts Valley, CA 95067-6989
Tel: 831-438-1565
E-mail: mark@pigeonpoint.com
Website: www.pigeonpoint.com