


## Platform management for PICMG architectures

By Mark Overgaard



**O**ne key challenge in deploying open modular platforms such as those based on the AdvancedTCA, CompactPCI, and the upcoming CompactTCA specifications is management, including platform management. As a foundation management layer, platform management inventories the Field Replaceable Units (FRUs) in the system's shelves, monitors their basic health, and manages their power, cooling, and interconnect resources. To meet this challenge, PICMG is leveraging the widely used Intelligent Platform Management Interface (IPMI) infrastructure, adding advanced extensions that address open modular systems' special platform management needs.

PICMG 3.0, the AdvancedTCA specification, defines the necessary extensions to IPMI. With the revisions in the just-adopted Engineering Change Notice updating PICMG 3.0 Rev 1.0, the specification's shelf management section accounts for almost 160 pages. In addition to mapping existing IPMI constructs into AdvancedTCA, this section defines 25 new commands, six new FRU Information Data Structures (several quite complex) and two new sensor types.

PICMG 2.9, the CompactPCI System Management specification, maps IPMI to the CompactPCI architecture, but does not address the necessary extensions. As a result, competing and incompatible sets of OEM-specific extensions occupy the CompactPCI space. The CompactTCA specification now in development aims to bring a subset of the AdvancedTCA extensions into the CompactPCI space, eliminating the need for these incompatible OEM-specific extensions. Even before the completion of the CompactTCA specification, the AdvancedTCA model provides an excellent basis for managing CompactPCI shelves, especially in the area of outside-the-shelf compatible interface support, enabling organizations with both AdvancedTCA and CompactPCI (or later CompactTCA) shelves to manage them in a coordinated fashion.

The IPM Sentry Intelligent Platform Management products are a set of off-the-shelf building blocks for the management subsystems of AdvancedTCA and CompactTCA products. This article offers examples of how these building blocks enable AdvancedTCA and CompactTCA product developers to focus on their unique added values, without the task of developing and validating management subsystems from scratch. The article focuses primarily on AdvancedTCA, but many of the management concepts for that architecture relate to Eurocard-based PICMG form factors as well.

### Intelligent platform management in AdvancedTCA: An overview

Figure 1 shows the logical elements of an AdvancedTCA shelf and identifies relevant AdvancedTCA specifications and potential sites for incorporation of IPM Sentry products or other off-the-shelf management building blocks.

An AdvancedTCA Shelf Manager communicates inside the shelf with IPM Controllers, each of which is responsible for local management of one or more FRUs, such as boards, fan trays, or power entry modules. Management communication within a shelf occurs primarily over the Intelligent Platform Management Bus (IPMB), implemented on a dual-redundant basis as IPMB-0 in AdvancedTCA.

An overall System Manager (typically external to the shelf) coordinates the activities of multiple shelves. A System Manager typically communicates with each Shelf Manager over Ethernet or

possibly another transport that supports the Internet Protocol (IP), such as Infini-band.

If Figure 1 depicted a CompactTCA or PICMG 2.16 CompactPCI shelf, it would have a nonredundant IPMB-0. In addition, there might be radial board select and healthy signals going to each slot from the Shelf Manager. Figure 1 shows three levels of management:

1. Board
2. Shelf
3. System

The next two sections address the board and shelf levels in turn, with a third section covering System Manager Interface by which a System Manager accesses a collection of shelves.

### Management aspects of an AdvancedTCA board

Figure 2 shows the key backplane interfaces of an AdvancedTCA board, including:

- Dual redundant -48VDC power from the backplane. A DC/DC converter powers the IPM Controller and, when enabled, power for the rest of the board.
- Dual redundant IPMB-0, the primary in-shelf management link to the IPM Controller. An 8-bit (including parity) hardware address establishes the IPMB-0 address through which a board communicates.
- Fabric, Base, and Update Channel Interfaces establish point-to-point backplane interconnects in a particular board and slot. The IPM Controller, through the Electronic Keying (E-Keying) process, enables the appropriate subset of the ports on those interconnects.
- Synchronization Clock Interface and Metallic Test Bus, both of which are based on the backplane. A separate E-Keying process avoids conflicting uses of these buses within a shelf.

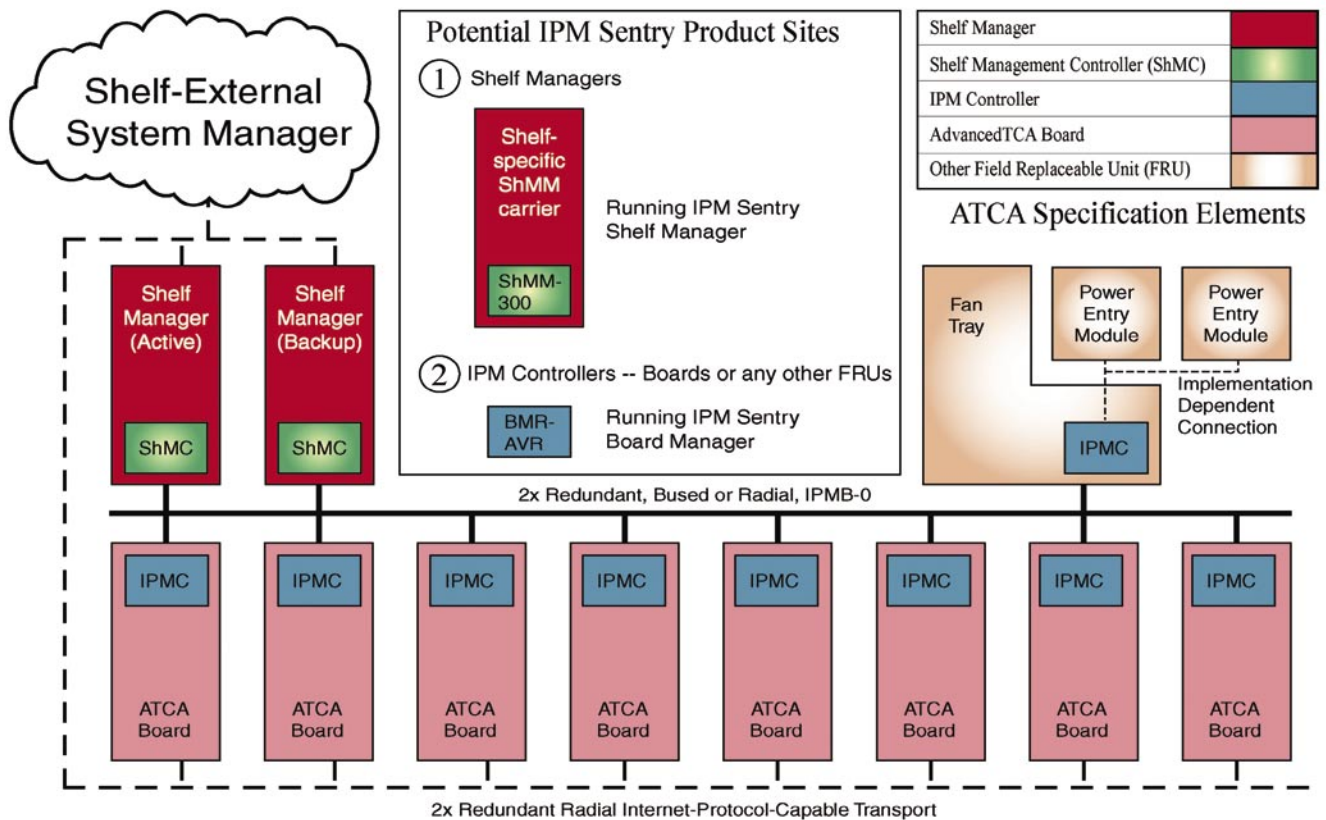


Figure 1

The *Payload* portion of an AdvancedTCA board implements the board's main functions and communicates with the IPM Controller over the Payload Interface. AdvancedTCA leaves the details of the Payload Interface entirely to the board implementer.

The IPM Controller uses non-volatile storage for various purposes, such as a repository for FRU information describing the Fabric, Base, and Update Channel Interfaces the board employs for use in E-Keying, discussed below.

AdvancedTCA requires a watchdog timer external to the controller in an IPM Controller. If this watchdog is ever allowed to expire, the IPM Controller is disconnected from IPMB-0 and undergoes a reset operation as a recovery measure. It is crucial that such a recovery reset not disturb the operation of the payload, and to meet this need many of the payload-oriented controls persist across such resets. Controls in this category include the Point-to-Point E-Keying and the Payload Power Enables.

In addition, AdvancedTCA boards must include at least one temperature sensor. When the board exceeds temperature thresholds, the Shelf Manager is notified to take corrective action.

As shown in Figure 2, a board's IPM Controller also manages key operator interfaces that include those listed below. Although not required by the specification, these operator interfaces are likely to be required by customers:

- The handle switch that is integrated with the lower board

handle to indicate whether the board handle is open or closed, enabling automatic responses to board insertion and extraction actions.

- The blue hot-swap LED that is fully on when the board is safe for extraction, fully off when extraction is not safe, and blinking during transitions between those two states.




Figure 2

- Some number of FRU LEDs that communicate overall status of the board and its various functions. AdvancedTCA defines interoperable mechanisms for management software to control LEDs, including blink rate, color, and *lamp test* functions.

Many of the functions described above apply to the IPM Controller subsystems of other AdvancedTCA FRUs, such as fan trays or power entry modules. Such modules don't need E-Keying support.

Furthermore, a subset of the above IPM Controller functionality also applies directly to the PICMG Eurocard architectures. The CompactTCA specification will likely reduce the mandatory features of an IPM Controller from AdvancedTCA to match reduced needs and maximize compatibility with existing PICMG 2.9 implementations.

## PICMG shelf management

The Shelf Manager watches over managed devices, reporting anomalous conditions to the System Manager and taking whatever corrective actions it can to prevent shelf failure.

Given data about the shelf's capabilities, especially its power capacities, the Shelf Manager negotiates power budgets with boards and other FRUs to maintain shelf operation within those capacities. Similarly, designers can configure sensors on boards or elsewhere in the shelf to monitor temperatures and issue event messages when temperatures go outside of prescribed bounds. In response to those events, the Shelf Manager can change fan levels or power budgets assigned to particular FRUs, even shutting them down if necessary.

Another important Shelf Manager function is the E-Keying mechanism. E-Keying ensures that boards only enable compatible interfaces over the point-to-point links that are implemented in a particular backplane between pairs of slots. This check for compatibility prevents incompatible devices from harming each other.

Each board's FRU information indicates the capabilities of each point-to-point port implemented in the Base, Fabric, and

Update Channel Interfaces. Additionally, the Shelf Manager accesses FRU information describing the connectivity implemented by the backplane. By coordinating the information from these two sources, the Shelf Manager decides which ports on each board should be enabled or disabled and conveys that information to each board's IPM Controller via IPMB-0. The IPM Controller takes the appropriate enable/disable actions with the ports. This interaction occurs when new boards enter the shelf and when users prepare boards for removal.

Sensor Data Records (SDRs), accessible from the device via IPMB-0, control access to each IPM Controller's resources. SDRs describe the sensors a given IPM Controller uses so that generic software can interpret them. For instance, temperature sensors can define threshold levels of increasing severity at specific temperature levels chosen by the designer, because such levels are likely inherently specific to the board's design.

## Shelf-external interfaces

The shelf-external interfaces provide access to the shelf from a System Manager, a logical concept that may include software as well as human operators in the *swivel chairs* of an operations center. Figure 3 shows the shelf-external interfaces of a typical AdvancedTCA Shelf Manager. The organization that specifies the interface and how fully it is specified distinguish the three interface categories:

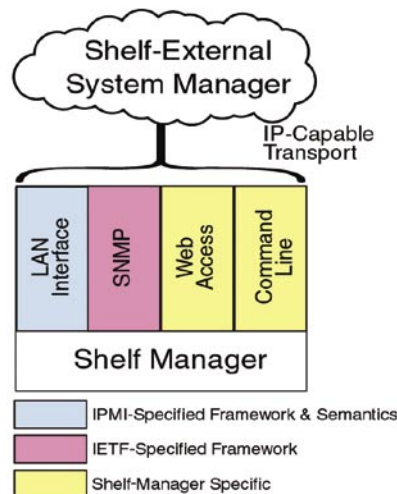


Figure 3

- IPMI LAN interface: IPMI specifies both the framework and the semantics for this interface, which is required for AdvancedTCA Shelf Managers.
- SNMP-based interfaces: For these interfaces, the framework is specified by the IETF, but (with one exception) the semantics are Shelf-Manager specific. IPMI Platform Event Traps that are SNMP trap messages with IPMI-defined semantics form the exception.
- Command Line and Web-based interfaces: There is no relevant open specification; these interfaces are necessarily Shelf-Manager specific.

AdvancedTCA requires an Internet Protocol (IP) capable transport for accessing the Shelf Manager. Typically Ethernet serves this purpose. Command Line Interfaces (CLIs), if available on a given Shelf Manager, are usually accessible either via Telnet over IP or through a direct serial connection.

## IPMI LAN interface

To ensure at least one interoperable means for the System Manager to communicate with a Shelf Manager, AdvancedTCA requires support for the IPMI LAN interface. This interface is very IPMI oriented (as Figure 3 shows), and an IPMI-aware System Manager can get visibility into the shelf through this interface for practically any significant IPMI aspect of the shelf. Furthermore, such a System Manager has equivalent visibility into any compliant Shelf Manager, facilitating the management of independently developed shelf products. Even better, the IPMI LAN interface can easily provide essentially equivalent visibility to either an AdvancedTCA- or Eurocard-based PICMG shelf.

## SNMP-based interfaces

Communication networks management uses SNMP extensively. For SNMP, the IETF specifies:

- The protocol by which management stations (in this instance, System Managers) communicate with management agents on managed nodes (in this instance, Shelf Managers representing their shelves).

- The framework for defining and navigating the Management Information Base (MIB) for accessing the management data on a managed node.

There is no generic specification-defined MIB for AdvancedTCA shelves. Furthermore, AdvancedTCA makes no requirements regarding SNMP support. Nevertheless, AdvancedTCA Shelf Managers will almost certainly include SNMP support (including a Shelf-Manager-specific MIB), and that support can compatibly cover CompactPCI/TCA shelves as well.

### Command Line and Web interfaces

Every AdvancedTCA Shelf Manager is likely to provide a CLI although the details will vary across implementations. The CLI is likely to be accessible by a direct serial connection to the Shelf Manager, or by an IP-capable transport such as Ethernet using Telnet. The IPM Sentry Shelf Manager treats the CLI as a connection to a specific Shelf Manager instance so that direct simultaneous CLI sessions can occur with both the active and backup Shelf Managers.

CLI use typically happens when a shelf is first installed. For instance, the AdvancedTCA Shelf FRU information defines two fields in the power distribution record that are intended to be configured to reflect actual power availability at the shelf's installation site. These fields describe the maximum current and minimum voltage supplied by the battery plant. The Shelf Manager bases its power availability estimates on these values and other data from the Shelf FRU information about the shelf's internal current carrying and distribution capabilities.

Any available Web interface to the Shelf Manager is also Shelf-Manager specific. Typically a subset of CLI-supported functionality, the Web interface gives simplified access to shelf functions. As with the other shelf external interfaces, it is straightforward for these interfaces to provide compatible AdvancedTCA and CompactTCA support.

### Using off-the-shelf management building blocks

Developing management subsystems for AdvancedTCA shelf- or board-level management (or doing the same for analogs in the CompactPCI variants) is a serious engineering project, even if you begin with a component that already complies with the IPMI specifications.

In addition to specification compliance, there is also the matter of achieving full interoperability with independently implemented components. To address this issue, PICMG is sponsoring a series of AdvancedTCA Interoperability Workshops (AIWs) where developers bring independently implemented product components (including a wide range of board or shelf products) together for systematic interoperability testing. Six such workshops have been held so far, with more planned.

Organizations developing new products that support platform management should definitely consider using off-the-shelf management subsystems. In choosing suppliers for these building blocks, look for demonstrated expertise in the relevant PICMG specifications and extensive interoperability experience.

### IPM Sentry: Off-the-shelf building blocks for PICMG platforms

This family of building blocks provides one way for developers of PICMG-compliant managed products to focus their efforts on their unique value-adds and leverage the work of management specialists.

A shelf developer can create a shelf-specific carrier for the IPM Sentry ShMM-300 shelf management mezzanine shown in Figure 4. At least one ShMM-300 carrier plugs into either an AdvancedTCA shelf or a CompactPCI shelf and automatically adapts to the differences. The ShMM-300 is a SODIMM-

sized (67.60mm x 50.80mm) module that includes a processor, Flash, and SDRAM memory. It is delivered with the IPM Sentry Shelf Manager software already installed in Flash memory.

This compact mezzanine approach allows great freedom in the physical placement of the Shelf Manager function, letting designers choose the carrier's physical form factor to match the planned placement. Functional side flexibility comes into play also. For instance, some shelf designers prefer to have the Shelf Manager directly manage the fans; others prefer that intelligent fan trays handle that mission. The ShMM-300 carrier can support either approach. Developers can license the IPM Sentry Shelf Manager's source code for even more flexibility.

At the level of a board (or other intelligent FRU), the IPM Sentry BMR-AVR reference design defines an IPM Controller that is ready to be integrated with the schematic of the rest of the board. In source form the IPM Sentry Board Manager software works with that hardware to create a compliant, interoperability-tested IPM Controller. The BMR-AVR design incorporates a pair of highly integrated Atmel AVR microcontrollers. The core of the IPM Controller core fits within a 18mm x 23mm area, so very little board real estate needs to be invested in the management subsystem. Figure 5 shows the physical placement of the main components of this core and area it occupies.

As of this writing, ten or more configurations of the IPM Sentry Shelf Manager




Figure 4




Figure 5

are shipping or in development for AdvancedTCA or CompactPCI. Close to twenty boards or other intelligent FRUs using the IPM Sentry board management solution are shipping or in development. Many of these have already been successfully interoperability tested at PICMG AIW events.

*Mark Overgaard* founded Pigeon Point Systems (PPS) in 1997 to focus on products and services supporting the adoption of open modular platforms to replace proprietary architectures, with an initial focus on the telecommunications market and CompactPCI. He is a leader in the technical subcommittees of PICMG (including the management aspects of AdvancedTCA and the corresponding CompactTCA specification, now in development). The current PPS product focus is the IPM Sentry line of platform management components, including AdvancedTCA shelf and board-level management components. Previously Mark was vice president of engineering at Lynx Real-Time Systems (a UNIX-compatible RTOS supplier) and TeleSoft (a major supplier of embedded development solutions for Ada). He earned an M.S. in Computer Science from UC San Diego and a B.S. in Physics from Geneva College.

For further information, contact Mark at:

**Pigeon Point Systems**  
P.O. Box 66989  
Scotts Valley, CA 95067-6989  
Tel: 831-438-1565  
E-mail: [mark@pigeonpoint.com](mailto:mark@pigeonpoint.com)  
Web site: [www.pigeonpoint.com](http://www.pigeonpoint.com)