Management building blocks speed AdvancedTCA shelf and system development

By Mark Overgaard

CompactPCI Systems

AdvancedTCA (PICMG 3.0) shelf management defines a platform management layer that inventories field replaceable units (FRUs) in the shelves of a system, monitors their basic health, and manages their power, cooling, and interconnect resources. AdvancedTCA leverages the widely used Intelligent Platform Management Interface (IPMI) infrastructure, adding advanced extensions that address the special platform management needs of communication systems.

Pigeon Point Systems' IPM Sentry shelf manager and ShMM-300 shelf management mezzanine products provide compliant and interoperability-tested building blocks to speed development of AdvancedTCA shelves. In addition, these building blocks implement a rich set of shelf external interfaces to facilitate the management of the one or more shelves that form a complete system. These interfaces bring building block benefits to the creation of the higher level System Manager software that coordinates an overall system, possibly comprised of multiple shelves.

In Integrating Shelf Management Solutions into CompactPCI and AdvancedTCA Chassis Designs, *Elma Electronic describes the use of IPM Sentry building blocks in AdvancedTCA and CompactPCI shelves.* Management building blocks speed AdvancedTCA product development *in the March 2003 issue of CompactPCI Systems provides more background on AdvancedTCA shelf management and the IPM Sentry products.*

Overview of AdvancedTCA shelf-external interfaces

In Figure 1, the shelf external interfaces of a typical AdvancedTCA shelf manager are shown. There are three categories of interface that are distinguished by the organization that specifies the interface:

- With IPMI LAN Interface, IPMI LAN specifies both the framework and semantics for this interface. Specifying these components of the interface is required for AdvancedTCA shelf managers.
- With SNMP-based interfaces the framework is specified by the IETF, but (with one exception) the semantics are shelf manager specific. The exception is IPMI Platform Event Traps. These are SNMP trap messages with IPMI defined semantics.
- Command line and web-based interfaces have no relevant open specification; these interfaces are necessarily shelf manager specific.

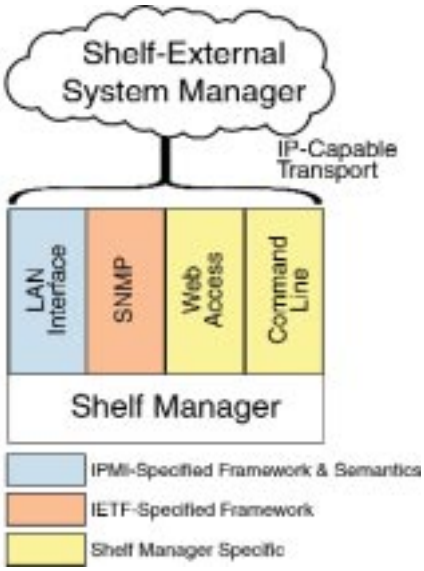

Figure 1

AdvancedTCA requires an Internet Protocol (IP) capable transport for accessing the shelf manager that is typically Ethernet. Command line interfaces, if available on a given shelf manager, are typically accessible either via Telnet over IP or through a direct serial connection.

The following sections cover each of these three categories of shelf external interfaces: IPMI LAN, SNMP-based, and Web-based.

IPMI LAN interface

AdvancedTCA requires that the IPMI LAN interface be supported to ensure that there is at least one interoperable means for the system manager to communicate with a shelf manager. In Figure 2, an IPMI aware system manager can get visibility into the shelf through the IPMI LAN 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.

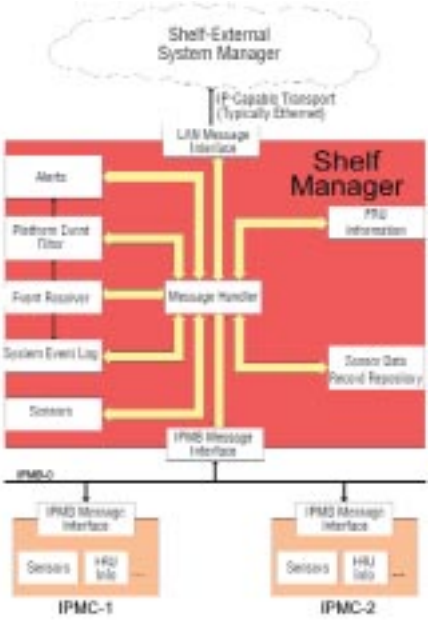

Figure 2

A key component in realizing the IPMI LAN interface is the message handler. For the LAN Interface, the message handler provides session and authentication services to improve the security and convenience of communication with the shelf manager. The message handler interfaces with the logical subsystems of the shelf manager to provide system manager visibility in areas such as the following:

- FRU information. This provides inventory information about the field replaceable unit that implements the shelf manager along with similar information about the overall shelf.

Shelf level information, for instance, includes the number of slots and the point-to-point connectivity implemented by the backplane in the base, fabric, and update channel interfaces.

- Sensor Data Record Repository. In AdvancedTCA, the sensor data record repository contains records describing all the IPMCs in the shelf and all the FRUs represented by those IPMCs.
- Sensors (those associated with the shelf manager itself). These sensors would include the state of the shelf manager's connection to IPMB-0 and other information.
- System Event Log. The system event log is a time-stamped nonvolatile record of the most recent platform events that have been logged in the shelf. Platform event results, for instance, when temperature sensors anywhere in the shelf exceed their programmed thresholds. The current state of the system event log is available to the system manager at any time. Platform events are fielded in the shelf manager by an event receiver function that dispatches them to both the system event log and the platform event filter.
- Platform Event Filter. The platform event filter function compares an event against a set of configurable event filters and triggers a selectable action on a match. One such action is the generation of an alert.
- Alert. The alert function implements alert actions triggered by the platform event filter. For the LAN interface, the primary type of alert is the platform event trap, which is an SNMP trap message populated with data about the event, sent to any of up to 15 configurable IP addresses. This facility could be used to have an SNMP trap message sent any time the FRU population in a shelf changes.

In addition to interacting with shelf manager maintained data and functions as described above, the system manager can also issue essentially arbitrary IPMI commands to IPMCs in the shelf. For instance, the system manager accesses sensors and FRU information of the IPMCs, or finds out which of the point-to-point interfaces have been enabled.

SNMP-based interface

A widely used interface for management of communications networks is SNMP. The IETF specifies the protocol that management stations (system managers in this case) communicate with management agents on managed nodes, and IETF specifies the framework for defining and navigating the management information base (MIB) that the management data on the managed node is accessed.

AdvancedTCA is without a generic specification or definition of a MIB for shelves. Along this same note, AdvancedTCA has a lack of requirements regarding SNMP support. However, AdvancedTCA shelf managers will almost certainly include SNMP support including a shelf manager specified MIB.

The IPM Sentry shelf manager provides SNMP access to almost every data and control function supported by the shelf manager. This includes the following groups of MIB variables:

- IPM controllers
- FRU devices
- Sensors
- Boards
- Shelf/shelves
- System event log
- LAN configuration parameters
- PEF configuration parameters

These groups of variables can be traversed to get detailed information or to exercise fine-grained control of the shelf and its FRUs, including full IPMI data.

Command Line and Web-based Interface

Every AdvancedTCA shelf manager is likely to provide a command line interface (CLI), however the details of the CLI will vary across implementations. The CLI is likely to be accessible by a direct serial connection to the shelf manager or via an IP capable transport such as Ethernet, using Telnet. In the IPM Sentry shelf manager, the CLI is considered a connection to a specific shelf manager instance so that direct simultaneous CLI sessions can be under way with both the active and backup shelf managers.

One likely time for the CLI to be used is 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 installation site of shelf. 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, the shelf's internal current carrying and current distribution capabilities.

A simple CLI dialogue with the IPM Sentry shelf manager is shown in Listing 1. There are four steps to this typical dialog:

1) Gather the current values of the power distribution record
2) Change the maximum available current for feed #1 to 8
3) Change the minimum voltage to -60
4) Show the modified values of the record after the shelf is restarted

For simplicity, this example assumes that the shelf is configured for a single power feed and a single slot.

The web-based interface provides straightforward access for much of the command line functionality from any HTTP capable browser. Unlike the CLI, the web-based interface is considered a connection to the active shelf manager and this connection automatically switches to the backup shelf manager if active.

Interfacing to Redundant Shelf Managers

AdvancedTCA recommends that each shelf support dual redundant shelf managers. In fact, a non-redundant shelf manager would represent a very serious single point of failure for a shelf. AdvancedTCA further requires support for a *shelf manager IP Address*, so the system manager can communicate with the shelf manager using an IP capable transport. In dual redundant configurations, this communication happens with the active shelf manager representing the shelf. If a backup shelf manager becomes active, this shelf manager takes over the role of representing the shelf using the IP address.

Shelf manager redundancy is an additional significant area in which AdvancedTCA extends IPMI, the later of which does not address the concept at all. By conscious choice, AdvancedTCA requirements focus on functional properties of the redundant pair, not on specifying details of how the members of the pair coordinate their actions. For instance, in the IPMI LAN interface, the specification requires that the redundant instances contain equivalent IPMI configuration parameters and FRU information at all times, except when the

```
**Step 1: Gather initial content of power distribution record.
CLI> shelf pd
Power Distribution:
Feed count: 1
Feed 00:
Maximum External Available Current: 0.0 Amps
Maximum Internal Current: 10.0 Amps
Minimum Expected Operating Voltage: -36.0 Volts
Actual Power Available: 0.000 Watts
Currently Used Power: 0.000 Watts
Feed-to-FRU Mapping entries count: 1
FRU Addr: 41, FRU ID: fe
**Step 2: Modify Maximum External Available Current.
CLI> shelf maxcurrent 0 8
Updating Shelf FRU Info
**Step 3: Modify Minimum Expected Operating Voltage.
CLI> shelf minvoltage 0 –60
Updating Shelf FRU Info
**Step 4: Check for updated record fields after shelf restart.
CLI> shelf pd
Power Distribution:
Feed count: 1
Feed 00:
Maximum External Available Current: 8.0 Amps
Maximum Internal Current: 8.0 Amps
Minimum Expected Operating Voltage: -60.0 Volts
Actual Power Available: 480.000 Watts
Currently Used Power: 10.000 Watts
Feed-to-FRU Mapping entries count: 1
FRU Addr: 41, FRU ID: fe
CLI>
*End pseudo code section
```

LISTING 1

active shelf manager is processing a command that modifies its data.

In Figure 3, one way that redundancy can be realized using IPM Sentry products is shown. With IPM Sentry, the active shelf manager is accessed via the shelf manager IP address and three protocols. When the backup shelf manager changes status to active, it will automatically take over interactions using the IP address, and the system manager components interacting with the shelf are essentially unaffected by the switchover. The three protocols are:

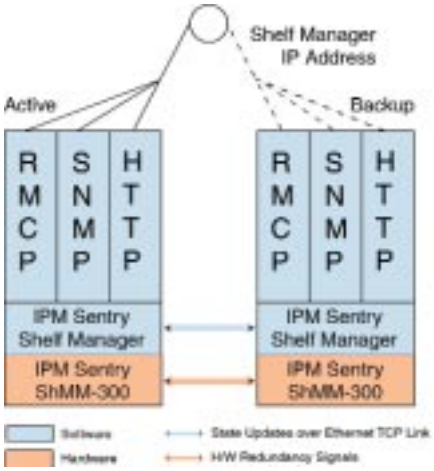

Figure 3

- Remote Management Control Protocol (RMCP). RMCP is used for the IPMI LAN Interface.
- SNMP, which provides an alternate access mechanism to essentially all functionality supported by the other interfaces.
- HTTP, which supports a subset of the command line interface from any Internet browser.

Mark Overgaard *is president of Pigeon Point Systems, which he founded 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 Software Interoperability chairman and most recently in the management aspects of AdvancedTCA).*

For more information, contact Mark at:

Mark Overgaard
Pigeon Point Systems
P.O. Box 66989
Scotts Valley, CA 95067
Tel: 831-438-1565
E-mail: mark@pigeonpoint.com
Web site: www.pigeonpoint.com