

An Introduction to the

Intelligent Platform Management Interface

The Intelligent Platform Management Interface (IPMI) is the standards-based systems management interface used by Dell™ PowerEdge™ servers. This article introduces the base specifications of IPMI technology and discusses revisions made to IPMI in its 1.5 and 2.0 versions.

BY JORDAN HARGRAVE

Providing consistent cross-platform systems management functionality has historically challenged IT administrators. The wide variety of hardware and operating systems initially required system vendors such as Dell to create proprietary systems management solutions, which usually involved Simple Network Management Protocol (SNMP) agents. These SNMP agents also required each system vendor to develop plug-ins for enterprise management consoles such as the HP® OpenView® or BMC Software® Patrol® applications. The disadvantage: vendors had no common model for providing access to general system information such as service tags, BIOS versions, or system type.

A history of management headaches

In the early 1990s, the Distributed Management Task Force (formerly the Desktop Management Task Force), or DMTF, attempted to remedy the cross-platform systems management problem. This organization—created by Dell, Intel, Compaq, and HP—aimed to develop software-based

specifications for systems management. Its first specification, the Desktop Management Interface (DMI), provided a large set of predefined object classes for a system, including hard disk, memory, CPU, and system chassis. DMI also included common objects for system events such as disk failure and chassis intrusion.

Any product that provided a standard DMI agent could be managed by a console that understood DMI. DMI 1.0 originally allowed querying only the local system; DMI 2.0 added the ability to query a system remotely. The next specification from the DMTF was the Common Information Model (CIM), which added object-oriented features, thereby allowing a greater reuse of code for common objects. This interface is currently used by Windows Management Instrumentation (WMI). The disadvantage of these purely software-based implementations, however, is that the server operating system (OS) must be up and running; if the system crashes, administrators cannot determine the root cause of the failure.

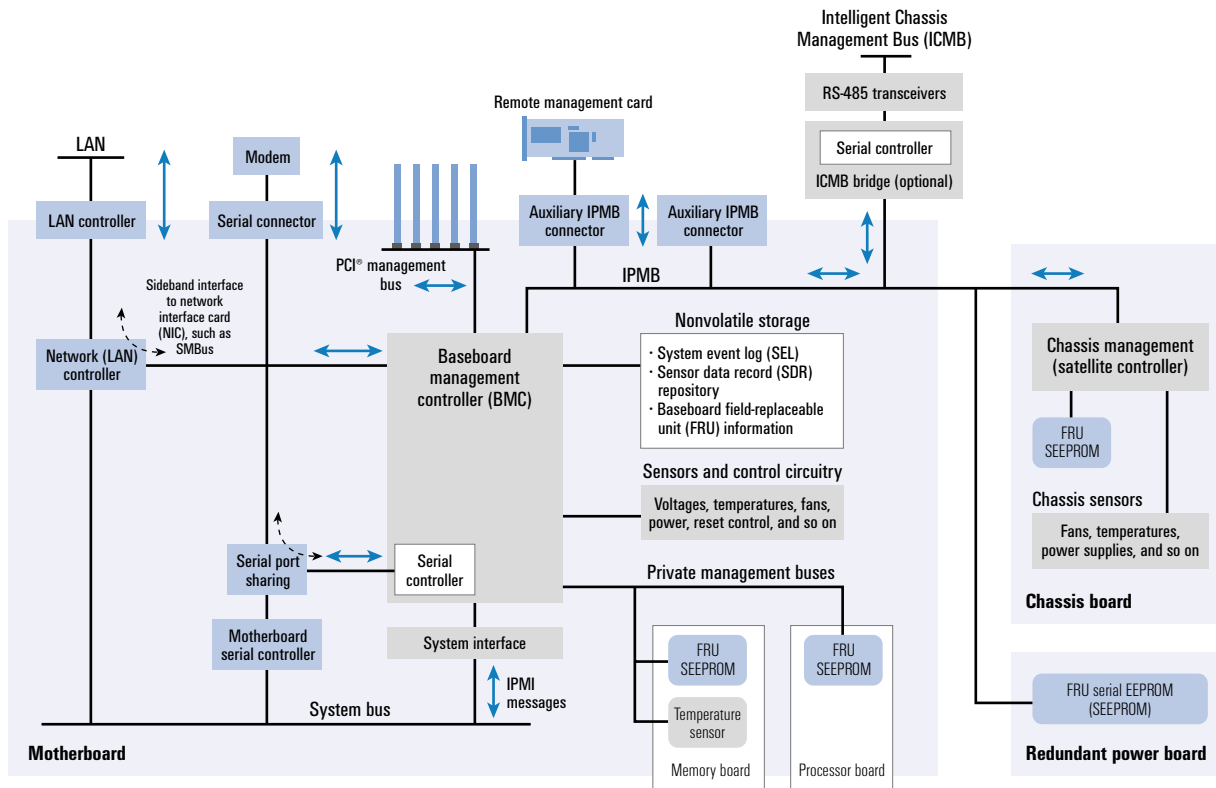


Figure 1. The IPMI architecture

Adding hardware to the mix

As systems became more powerful, they required increasingly robust remote systems management capability. This functionality was transferred from the software level to add-in or embedded controllers. Initial offerings from Dell, such as the Dell Remote Access Card (DRAC), were proprietary products that added new features to servers, such as a remote console and remote floppy boot. These cards can operate even if a server loses power or the OS crashes. They also can manage a server in a pre-OS state, such as during power-on self-test (POST) and BIOS configuration. Such hardware-based systems management implementations are complementary to the software agents. DMI/SNMP agents can provide most of the sensor and system information, while the card performs failure analysis at a remote console. The disadvantage of these cards is their proprietary nature—the agents and tools must be rewritten for each new release.

A new direction in systems management

In 1998, Intel, Dell, HP, and NEC developed the Intelligent Platform Management Interface (IPMI) as a specification for providing systems management capability in hardware. The IPMI specification provides a common message-based interface for accessing all the manageable features in a system. IPMI includes a rich set of

predefined commands and interfaces organized by type of operation, such as reading temperature, voltage, fan speed, and chassis intrusion. Methods are also provided for accessing the system event log (SEL), hardware watchdog, and power control.

IPMI replaces or abstracts previous methods of accessing sensors through the systems management bus (SMBus) or intelligent interface controller (i2c). The IPMI specification is expandable, allowing original equipment manufacturers (OEMs) to define their own commands or sensor types. Many companies such as AMD, Fujitsu, and QLogic have developed products that use IPMI. Several releases and enhancements have been made to IPMI since its original 0.9 specification; the latest release is 2.0.

Understanding how IPMI works

The Baseboard Management Controller (BMC) is the heart of an IPMI-based system—it is responsible for monitoring and controlling all the manageable devices in the system. Figure 1 illustrates the central role of the BMC within the IPMI architecture. All access from the host OS is routed through the BMC, which can talk to other IPMI-aware devices in the system through the Intelligent Platform Management Bus (IPMB). Several vendors have developed inexpensive embedded BMC solutions, providing a cost-effective method for administrators to add IPMI support to a system. These

controllers reduce server CPU usage by offloading system polling from the CPU.

The BMC contains nonvolatile RAM (NVRAM) storage for the SEL, sensor data records (SDRs), and asset information. The SDR area describes sensors that may be connected to the system. It stores information such as the sensor name, location, and thresholds. The BMC also is responsible for sending and handling events. These events can be thresholds that have been exceeded or triggers such as chassis intrusion. Actions to handle the events can include logging, power cycling, or issuing an SNMP trap (IPMI 1.5 only).

Several tools can be used for querying IPMI information on a system. The OpenIPMI project is an open source initiative to develop a suite of IPMI management utilities. It currently addresses only Linux® operating systems but may soon become available for other platforms. Intel has provided a set of reference IPMI drivers for Microsoft® Windows®, Novell® NetWare®, and UnixWare® operating systems.

Integrating IPMI into Dell PowerEdge servers

As one of the four original IPMI promoters, Dell is committed to using IPMI as its management interface for the Dell™ PowerEdge™ server line. Most of the current sixth-generation PowerEdge servers—the 1650, 2600, 2650, 6600, and 6650 models—support the IPMI 1.0 specification. The PowerEdge 1750, 3250, 7150, and 8450 also support IPMI. Dell OpenManage™ Server Assistant agents use IPMI to query server health, retrieve event logs, and perform power-control operations. The Dell Embedded Remote Assistant (ERA) and DRAC III cards communicate with a PowerEdge server through the IPMB. Dell plans to support IPMI 1.5 in future seventh- and eighth-generation PowerEdge servers.

The future of IPMI


Development of the IPMI specification is ongoing. The IPMI 1.5 specification, the latest approved specification, adds support for accessing the BMC through the serial port or a local area network (LAN). The LAN and serial connectors can be either dedicated to the BMC or multiplexed with the system connectors. This feature allows a server to be completely controlled from a remote system

through the network or a modem. The advantage of this feature is that no agents are required on the remote system. Existing tools that already use IPMI can be easily modified to support the new transport methods, because the message interface remains the same. The LAN transport uses the Remote Management Control Protocol (RMCP), which employs User Datagram Protocol (UDP) datagrams. The server also can automatically send event notifications either through the serial or LAN port as SNMP traps, or to a pager.

The IPMI 1.5 specification has added features for controlling the system boot order. This allows the remote system to boot to a utility partition or to boot through the Preboot Execution Environment (PXE). Another new feature is terminal mode, which provides a remote text-mode console that can be used for viewing the BIOS configuration screen or server crash dump.

Furthermore, the IPMI 2.0 specification is currently under review. This version adds enhanced features for authenticating and encrypting the LAN/serial connection.

A new framework for cross-platform systems management

IPMI is an important step in the evolution of server management. It is available in a wide range of server and system platforms and can help reduce total cost of ownership by providing a consistent interface across systems. The large set of common commands and the ability to manage a failed system remotely can provide powerful tools for IT administrators. 

Jordan Hargrave (jordan_hargrave@dell.com) is a senior software engineer at Dell. He has a bachelor's degree in Mathematics and Computer Science from Carnegie Mellon University.

FOR MORE INFORMATION

Intelligent Platform Management Interface:

<http://developer.intel.com/design/servers/ipmi>

OpenIPMI project:

<http://www.sourceforge.net/projects/openipmi>