

Remote Management with the Baseboard Management Controller in Eighth-Generation Dell PowerEdge Servers

The first wave of eighth-generation Dell™ servers—the Dell PowerEdge™ 1850 and PowerEdge 2850—is equipped with baseboard management controller (BMC) technology for enhanced flexibility in remote management. This article focuses on configuration of the BMC for remote management operations using the Dell OpenManage™ Server Administrator tool set.

BY HAIHONG ZHUO; JIANWEN YIN, PH.D.; AND ANIL V. RAO

The latest generation of Dell PowerEdge servers—including the recently released PowerEdge 1850 and PowerEdge 2850—is equipped with an on-board micro-controller called the baseboard management controller (BMC), which is compliant with the Intelligent Platform Management Interface (IPMI) 1.5 specification.¹ The IPMI specification defines standardized, abstracted interfaces to the platform management subsystem. The BMC is the heart of the IPMI architecture and provides the intelligence behind intelligent platform management—the autonomous monitoring and recovery features implemented directly in platform management hardware and firmware.

IPMI defines a set of common interfaces to the hardware and firmware that is used to monitor system health and manage the system. Because operating, administering, and maintaining servers can represent a significant portion of total cost of ownership (TCO), any reduction in these processes through automation can help reduce TCO. The BMC provides management capabilities that are designed to reduce TCO through asset tracking; automatic alerting,

shutdown, and restart; and manual remote restart and power control. Because the BMC is IPMI 1.5 compliant, IPMI reference documents present a thorough supplement to this article.²

One significant addition to the IPMI 1.5 specification is an interface that enables standard messaging and alerting through direct serial, IPMI over local area network (LAN), and serial over LAN (SOL) connections. This IPMI 1.5 feature helps provide administrators with remote access to manage the platform, even when the operating system (OS) or systems management software is not available or when the system is powered off. The standard also defines an alerting mechanism for the system to generate and send a Simple Network Management Protocol (SNMP)-based Platform Event Trap (PET) when a monitored event condition occurs. IPMI messaging and alerting enable remote management of the BMC: IPMI messaging can be used to query platform status, review hardware logs, or issue other requests from the remote console, which can also be notified of events or alerts through IPMI alerting functions.

¹For more information about the Intelligent Platform Management Interface, visit <http://www.intel.com/design/servers/ipmi/index.htm>.

²To download Intelligent Platform Management Interface reference documents, visit <http://www.intel.com/design/servers/ipmi/spec.htm>.

IPMI is best used with systems management software running under the OS, which offers an enhanced level of configuration and manageability by providing in-band (OS-present) access to IPMI management information and by integrating IPMI with additional management functions. The Dell OpenManage suite of systems management software has been enhanced to provide sophisticated management for the BMC that is compliant with IPMI 1.5.

As shown in Figure 1, five of the suite's components are related to the configuration and use of the BMC remote management capability:

- **Dell OpenManage Server Administrator:** This tool enables administrators to discover, configure, and manage the BMC, including settings for remote management. It provides both a graphical user interface (GUI) and a command-line interface (CLI).
- **Dell OpenManage Deployment Toolkit:** This tool enables administrators to configure the BMC in a pre-OS deployment phase. It provides a CLI to configure the BMC for remote management.
- **BIOS Setup (Option ROM) Utility:** This utility can be invoked from the BIOS power-on self-test (POST) sequence and enables administrators to configure a subset of the BMC for remote management. It sets up the most commonly used, quickest configuration of the BMC.
- **Dell OpenManage IT Assistant:** This management console can discover instrumented systems, as well as manage, receive, and resolve traps generated and sent by instrumentation and the BMC in the form of PETs.
- **BMC Management Utility:** This remote console program can be used to remotely access and manage the BMC, provided that the BMC is configured properly using Dell OpenManage Server Administrator. Administrators can use this utility through the CLI application ipmish.exe.

This article focuses on the configuration of the BMC for remote management using Dell OpenManage Server Administrator.

Configuring the BMC for remote access

To use the BMC Management Utility for remote management, administrators must configure attributes for the connection, which can be direct serial, IPMI over LAN, or SOL. This configuration can be performed using the Server Administrator GUI or CLI.³ The GUI can be accessed through a Web browser on a trusted domain, and the CLI can be run on the local server. Most features present in GUI have counterparts in the CLI for scripting convenience.

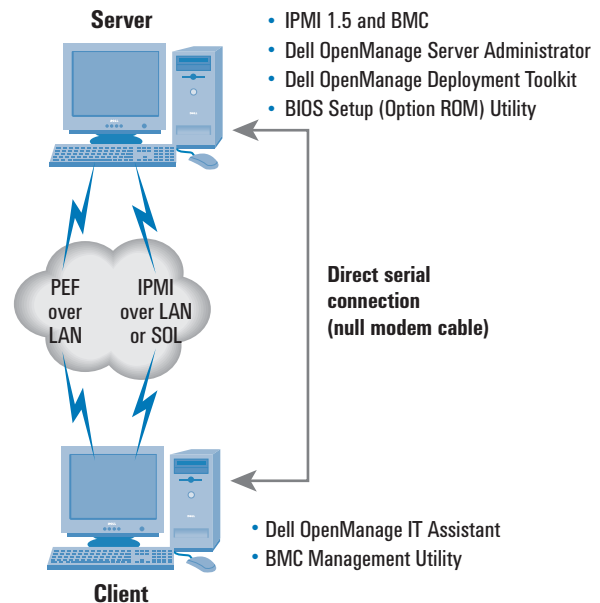


Figure 1. Remote management using the BMC

Direct serial connection

Over a direct serial connection, administrators can use a remote console to access PowerEdge servers equipped with an IPMI 1.5–compliant BMC. Because the connection is typically a point-to-point cabled connection, it can be configured in Basic mode, which uses a simple, clear-text password to activate a session. IPMI messages are encoded and delimited using a simple framing scheme based on escaped characters.

Another mechanism for BMC direct serial connection is Terminal mode, in which a simple VT-100 terminal or terminal emulator can be used to generate requests and receive responses from the BMC. Administrators can enter IPMI messages into the terminal using printable ASCII characters.

Attributes for a BMC direct serial connection can be configured through Server Administrator (see Figure 2) and include:

- **Connection mode setting:** Can be set to Direct connect Basic mode or Direct connect Terminal mode
- **IPMI messaging communication settings:** Includes baud rate and flow control
- **Channel privilege level limit:** Specifies the maximum privilege that a user can have on the serial connection

By selecting “Terminal Mode Settings” (circled in Figure 2), an administrator is shown another screen where settings for Terminal mode—including line editing, delete control, echo control,

³ For more information about the Dell OpenManage Server Administrator GUI and CLI, consult the *Dell OpenManage Server Administrator User's Guide*, available at <http://docs.us.dell.com/support/edocs/software/svradmin/index.htm>.

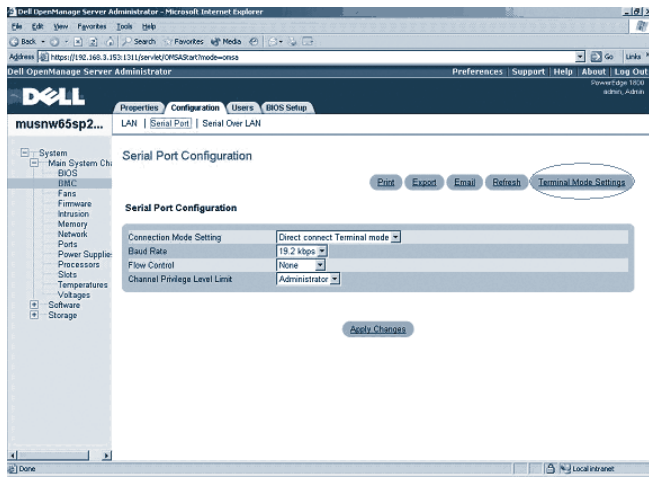


Figure 2. Using the Server Administrator GUI to configure the BMC serial connection

handshaking control, new line sequence, and input new line sequence—can be configured. Changing and applying any values from the GUI will save the new values in the BMC. The GUI is capability driven, so it displays only features that are supported on the particular server or IPMI implementation.

Any setting available through the Server Administrator GUI can be configured through the Server Administrator CLI. Because the CLI is not capability driven, it may offer more choices for some settings to support multiple platforms.

IPMI over LAN connection

The BMC on PowerEdge servers uses Remote Management and Control Protocol (RMCP) under IP version 4 to communicate with a remote console over an Ethernet LAN. This type of connection is referred to as IPMI over LAN. On the PowerEdge 1850 and PowerEdge 2850 servers, the IPMI over LAN interface is established using the embedded LAN controller—LAN on Motherboard 1 (LOM1)—that is shared by both the BMC and the system. The BMC has its own Media Access Control (MAC) address and IP address, which are different from the MAC address and IP address shown by the OS for LOM1.

Before IPMI messaging can work on a LAN connection, administrators must enable the IPMI over LAN mode. By default, the BMC is configured with this mode in the off position to prevent unauthorized access. However, even if IPMI over LAN is disabled, other related attributes can still be configured through Server Administrator and will take effect whenever IPMI over LAN is finally enabled. Parameters for the BMC LAN connection that can be configured through the Server Administrator GUI (see Figure 3) and CLI include:

- **IP address source:** Can be set to static or Dynamic Host Configuration Protocol (DHCP)
- **IP address settings:** Are available if the IP address source is static

- **Channel privilege level limit:** Specifies the maximum privilege that a user can have on a LAN connection
- **Virtual LAN (VLAN) tag:** Sets VLAN attributes

Server Administrator will validate the entered values by performing a range check and prompting administrators with a warning message if a value is out of range. Administrators can further explore these parameters by consulting the *Dell OpenManage Server Administrator User's Guide*, online help provided by the Server Administrator GUI, and the IPMI specifications.

SOL connection

SOL is the specification of packet formats and protocols for transmitting serial data over the LAN using IPMI over LAN packets. The typical goal of this capability is to redirect the traffic to and from a local asynchronous serial controller interface. This redirection enables communication over the LAN with local software that can only communicate through a local serial controller. Traffic to and from the local asynchronous serial controller interface on the baseboard can be routed to the BMC, SOL, or the system's serial port connector. To correctly route SOL traffic, administrators must select the "BMC NIC" option for the Serial Port 1 configuration of the Serial Configuration tab in the Server Administrator BIOS Setup page.

Like IPMI over LAN, SOL must be enabled before it can work. Attributes that must be configured for SOL to work properly include baud rate, channel privilege level limit (the maximum privilege that a user can have on the SOL connection), and advanced settings such as character accumulate interval and character send threshold.

Configuring BMC users for remote login

After a BMC remote management connection has been configured, users can access a BMC remotely through a console program, such

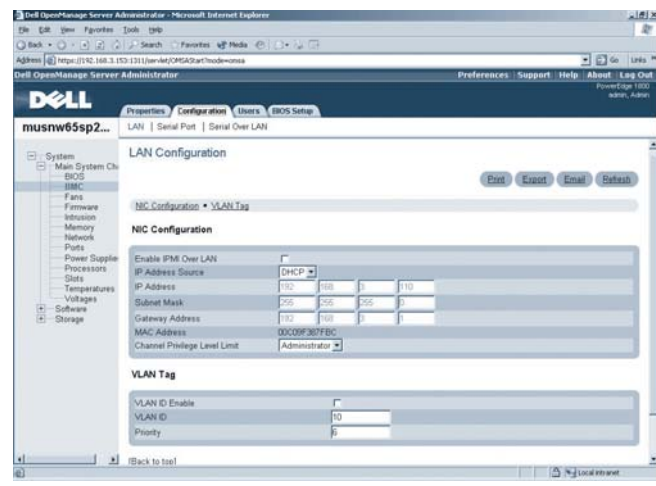


Figure 3. Using the Server Administrator GUI to configure the BMC LAN connection

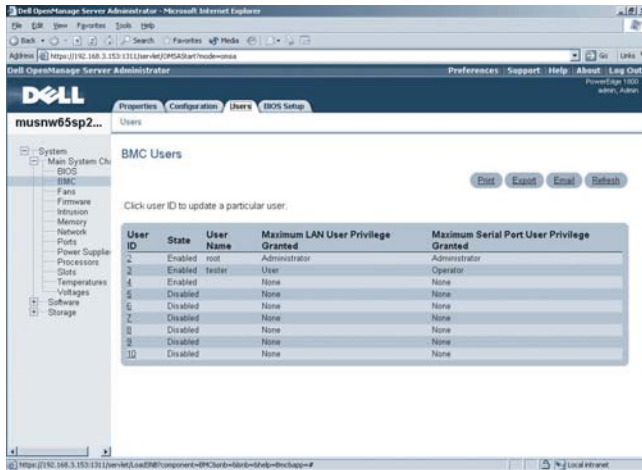


Figure 4. Status of BMC Remote Access Users

as the BMC Management Utility, to query the system's status or issue requests to perform actions. Users need to log in and authenticate to the BMC through the remote console program. The BMC maintains a local database of remote access users and their privileges. Each of these remote access users can be enabled for the serial connection, the LAN connections, or both, and they can have different privilege levels on each connection. Moreover, each connection has its own privilege level that can be configured and takes precedence over the privilege level of the user.

Figure 4 shows the BMC users that are configurable on PowerEdge 1850 and PowerEdge 2850 servers. Selecting the ID number of a user brings up a user configuration page, where administrators can update the user's state (enabled or disabled), name, password, and privilege level for each connection.

Per the IPMI 1.5 specification, User ID 1 is reserved for the anonymous user, a predefined user that has no user name or password. For security concerns, this user is disabled on Dell PowerEdge servers and therefore is not shown in the list of users.

Administrators can enable a user ID without assigning a user name, such as User ID 4 in Figure 4. However, when this user attempts to log in to the BMC from a remote console, BMC will try to authenticate the connection session with User ID 1, namely the anonymous user. Because the anonymous user is disabled on PowerEdge servers, the authentication will fail. Thus, a user ID without a user name, even if enabled, cannot be used to log in to the BMC remotely.


Configuring the BMC PET alert

When an event occurs on the platform, the BMC monitors this event, generates an event message, and logs the event in the BMC System Event Log (SEL). The BMC will, in turn, check whether the

event meets some of the configured Platform Event Filtering (PEF) criteria. If yes, the BMC generates an SNMP PET alert and sends it to the designated destination(s) over the LOM shared with the system. On PowerEdge 1850 and PowerEdge 2850 servers, because the BMC has its own IP address, the source of the SNMP trap will be the IP address of the BMC or the Domain Name System (DNS) hostname assigned to the BMC IP address, instead of the IP address shown in the OS for the LOM of the failing system. To provide information for SNMP manager applications to associate a PEF alert with the failing system, the NetBIOS hostname of the failing system is embedded in an original equipment manufacturer (OEM) variable binding. Dell OpenManage IT Assistant extracts this value and includes it in the description of this trap as extra event information.

For platform events to be filtered properly and the PET to be generated and sent out as desired, administrators must configure a platform event filter table, which contains all configured platform event filters, and alert actions associated with each entry in the table. Administrators should also configure the destinations and SNMP community to which the PETs are sent. The destination system should have an SNMP manager program, such as IT Assistant, that can receive, recognize, filter, and act on PETs.

Enabling better systems management through BMC and IPMI

Remote management of server platforms has historically been offered through add-on proprietary controllers and solutions. However, eighth-generation Dell servers such as the PowerEdge 1850 and the PowerEdge 2850 help to provide this level of manageability in the base platform through a BMC that employs industry-standard IPMI interfaces. This baseline remote management capability can help pave the way for the development of complex solution stacks that can be more easily managed. 

Haihong Zhuo is a software engineer in the Dell Enterprise Software Development Group. She has worked on systems management solutions and is currently on the Systems Management Instrumentation team. Haihong has an M.S. in Computer Engineering from The University of Texas at Austin and a B.S. in Electrical Engineering from Tsinghua University.

Jianwen Yin, Ph.D., is a software engineer senior consultant on the Managed Node Development team in the Dell Product Group. Currently, he is developing version 2.0 of Dell OpenManage Server Administrator. Jianwen has a Ph.D. in Computer Science from Texas A&M University.

Anil V. Rao is a senior software engineer and strategist in the Dell OpenManage Group, where he focuses on Dell OpenManage Server Administrator and systems management. At Dell, he has worked on RAID design and development as well as embedded, remote, and systems management software solutions. Anil has an M.S. in Computer Science from Texas Tech University and a B.S. in Computer Engineering from the University of Bombay.