

Troubleshooting a Modular Computing Environment

Modular computing environments offer several benefits such as easier deployment. The Dell™ OpenManage™ systems management software suite, which includes IT Assistant and Server Administrator, helps administrators monitor and manage components within Dell PowerEdge™ 1655MC server blade systems. This article describes how administrators can recognize, locate, and resolve errors in modular environments and improve administrative efficiency.

BY HAIHONG ZHUO AND WEIMIN PAN

Modular architectures segment existing servers into more manageable pieces. A modular server system contains several modules, including server blades, power supplies, Ethernet switches, and systems management modules. These modules often contain components of their own; for example, a server blade consists of a circuit board containing processors, memory, hard disks, and network connections. Server blades stack side by side in a common enclosure and share power, cooling hardware, and cabling. This type of modular architecture decreases rack space requirements, lowers power consumption, reduces the number of cables, lowers costs, improves manageability, and makes service easier. Several modular server systems compose a modular environment.

Looking inside the Dell PowerEdge 1655MC server blade system

As Figures 1 and 2 illustrate, one Dell™ PowerEdge™ 1655MC server blade system can contain up to six server blades, which share two power supplies, four cooling fans, two integrated Gigabit Ethernet¹ switches, and one

management module called the Embedded Remote Access/MC (ERA/MC) controller. Some modules, such as server blades, integrated switches, and the ERA/MC, can contain management software or firmware. These instrumented modules can be managed by higher level systems management software.

Other modules, such as power supplies and cooling fans, do not have management software or firmware installed. These uninstrumented modules, also known as enclosure infrastructures, are monitored and managed through the ERA/MC. Administrators also must manage individual components within some modules, such as the processors, memory, temperature probes, and disk storage devices of the server blades. The Dell OpenManage™ software suite provides the tools necessary for administrators to manage the PowerEdge 1655MC server blade system and its modules.

Understanding error messages from different modules

Within a modular server system, errors can occur on various modules, such as a memory error-correcting

¹Gigabit Ethernet indicates compliance with IEEE® 802.3ab and does not connote speeds of 1 Gbps.

code (ECC) error on a server blade, or a power supply failure in the enclosure.

Errors from server blades

Within a server blade, errors can occur in processors, memory, temperature probes of the processors, and disk storage devices. When an error occurs, the health status of the failing component changes from normal to warning or critical, and so does the health status of the server blade containing that component. The LED display on the front cover of the blade will generally change color, start blinking, or both, to indicate an error in that server blade.

Administrators can use Dell OpenManage Server Administrator to monitor and manage server blades. Server Administrator reports errors in the system event log and sends Simple Network Management Protocol (SNMP) traps to any designated managers.

Figure 3 lists some possible server blade failure conditions, the resulting health status of each event, how each status change is reported in the event log, and the corresponding SNMP trap message.

Errors from integrated Ethernet switches

Each integrated Ethernet switch has firmware installed that will perform a power-on self test (POST) check when the switch is powered up along with the enclosure. If the check fails, the switch initialization will hang, and the Fault Indicator LED on the switch will change to amber, indicating that the switch is not working properly. The ERA/MC will then detect this error and update its own status accordingly.

On the other hand, if the switch is working properly, a link-up or link-down event will be recorded in the local event log. Because the firmware supports SNMP, an informative trap will be sent to any designated managers. Figure 4 lists some errors that could occur on an integrated Ethernet switch.




Figure 1. The Dell PowerEdge 1655MC server blade system—front and back views

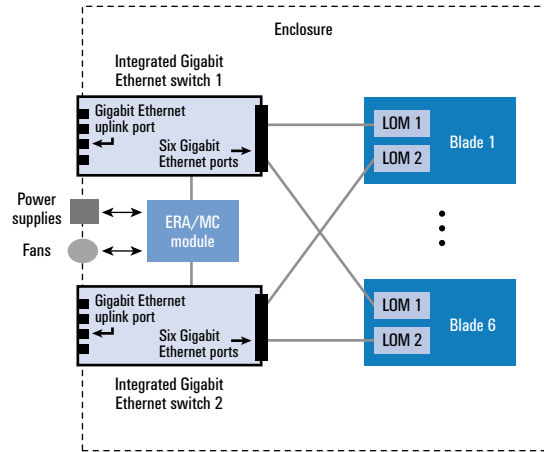


Figure 2. Logical view of the Dell PowerEdge 1655MC server blade system

Errors from the shared enclosure infrastructure

Because the shared enclosure infrastructure components, such as power supplies and cooling fans, are monitored and controlled by the ERA/MC, a change in the health status of these components alters the ERA/MC status as well. If an error in the shared enclosure infrastructure causes a change from a normal condition to warning or critical, the failed module’s Fault Indicator LED indicates an abnormal state and the ERA/MC status indicator LED changes to amber. The ERA/MC event log also records the error. The ERA/MC has an embedded SNMP agent, which sends an SNMP trap to any designated manager when an error occurs in a shared enclosure module or one of its components. Figure 5 lists some possible shared enclosure module errors.

Failure condition	Status	Event log	SNMP trap
DIMM ECC error	Critical	Memory device status is critical	Memory device status is critical [Memory device location: DIMM_A]
Disk drive failure	Critical	Drive slot sensor drive error	Drive slot sensor drive error
CPU temperature over warning threshold	Warning	Temperature sensor (CPU1) detected a warning	Temperature sensor detected a warning
CPU temperature over critical threshold	Critical	Temperature sensor (CPU1) detected a failure	Temperature sensor detected a failure

Figure 3. Examples of server blade errors

Failure condition	Status	Event log	SNMP trap
Port link down	Normal	Port [n] link down	Port [n] link down
POST failure	Switch hang	N/A	N/A

Figure 4. Examples of integrated Ethernet switch errors

Failure condition	Status	Event log	SNMP trap
Power supply DC failure	Warning	Power supply sensor failure detected	Power supply sensor failure detected
Power supply AC failure	Warning	Power supply sensor power lost	Power supply sensor power lost
Power supply removed	Warning	Power supply removed	Power supply removed
Fan speed below/above warning threshold	Critical if more than two fans, warning if one fan	RPM fan sensor detected a warning	RPM fan sensor detected a warning (x RPM) [where x represents number of RPMs]
Fan speed below/above critical threshold	Critical if more than two fans, warning if one fan	RPM fan sensor detected a failure	RPM fan sensor detected a failure (x RPM) [where x represents number of RPMs]
Fan removed	Critical if more than two fans, warning if one fan	RPM fan sensor detected a non-recoverable event (0 RPM)	RPM fan sensor detected a non-recoverable event (0 RPM)
Enclosure temperature exceeded warning threshold	Warning	Temperature sensor detected a warning	Temperature sensor detected a warning
Enclosure temperature exceeded critical threshold	Critical	Temperature sensor detected a failure	Temperature sensor detected a failure

Figure 5. Examples of shared enclosure module errors

Identifying the failing module

The centralized management console of the Dell OpenManage IT Assistant software² offers a starting point for troubleshooting the PowerEdge 1655MC server blade system. IT Assistant discovers the instrumented modules of the PowerEdge 1655MC—server blades, integrated switches, and the ERA/MC—and automatically places them in a custom group that represents the enclosure. Figure 6 shows an example of a custom group named SHREDDER_2; IT Assistant derives the group name from the service tag of the enclosure, but administrators can rename a group to something more meaningful.

In the system tree, IT Assistant lists modules by their host names: for example, itabld21, itabld22, and itabld23 for server blades (see Figure 6). If no host name was ever assigned to an integrated switch, a default host name will be assigned when the switch is powered up. The default name follows the format SW_n_CHASSIS-SERVICE-TAG where *n* is the number of the switch (1 for the top switch, 2 for the bottom switch) and CHASSIS-SERVICE-TAG is the service tag of that server blade system’s enclosure, or chassis. For example, if the service tag is SHRDIT2, the two integrated switches will be listed as SW1_SHRDIT2 and SW2_SHRDIT2 (see Figure 6). The same rule applies to the ERA/MC, where the format is RAC_CHASSIS-SERVICE-TAG. In

Figure 6, the name RAC_SHRDIT2 represents the ERA/MC belonging to the same enclosure as the integrated switches.

Figure 6 shows that each module and its components, if any, are in healthy status—noted by a green check mark in the IT Assistant system tree.

Troubleshooting server blades

Dell OpenManage offers several methods for troubleshooting problems within server blades. These methods include checking health status, viewing SNMP traps and event logs, running diagnostics, and using IT Assistant to launch the ERA/MC Web console and physically locate a failed server blade.

Detecting health status

When the health status of a component within a server blade changes, so does the health status of the server blade itself. IT Assistant detects these status changes if the server blade is instrumented with OpenManage Server Administrator. For example, if a memory error occurs in a server blade, as illustrated in Figure 7, IT Assistant shows the memory (DIMM_B) in a critical state.

Troubleshooting using Server Administrator. OpenManage Server Administrator can provide additional information for further investigating a problem on a server blade. This program can be launched from IT Assistant by selecting Server Administrator from the Status page of the server blade and then clicking the Server Administrator button. A new Web browser window that displays the Health page for that server blade will appear.

In the scenario where DIMM_B is in a critical state, the Server Administrator Health page will show a red “X” in the Severity column of the Main System Chassis object. Clicking on “Main System Chassis” displays the status of the monitored components. In this scenario, the Memory component will also have a red “X”

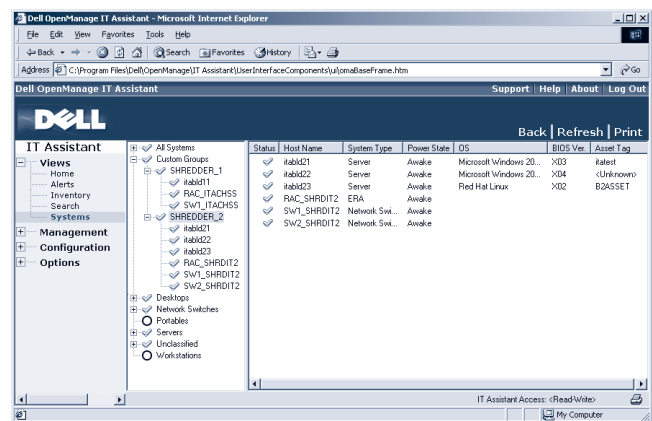


Figure 6. Server blades discovered by IT Assistant

²For more information about IT Assistant, see “Using IT Assistant to Manage Blade Servers” by Manoj Gujarathi in Dell *Power Solutions*, August 2002.

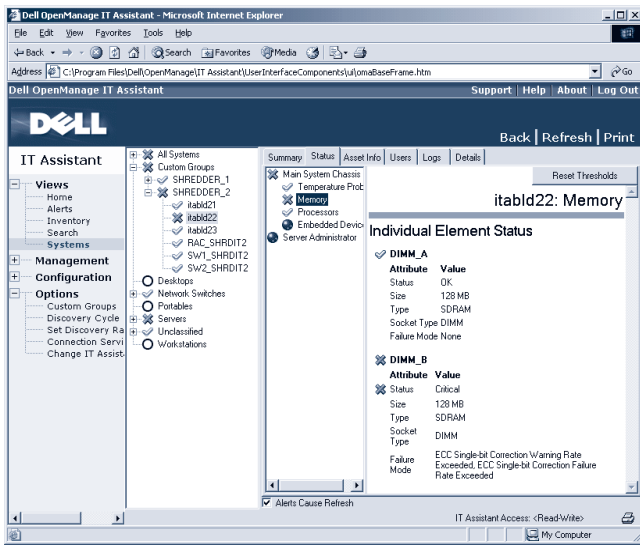


Figure 7. Memory failure in a server blade

as its status indicator. Clicking on “Memory” brings up the Properties tab, which shows that DIMM_B is in a critical state. The administrator can click the DIMM_B object for detailed memory device and failure information.

Using traps and event logs

OpenManage Server Administrator, when installed on a server blade, will send an SNMP trap to inform IT Assistant of any events if the trap community and destination are set correctly.

To set the trap community and destination:

1. Open the Service Control Manager on the appropriate server blade
2. Double-click “SNMP Service” to modify SNMP properties
3. Select the Traps tab
4. Enter the appropriate community name in the Community name box
5. Add the destination IP address or host name into the trap destination list

The Alerts page of IT Assistant displays traps received by IT Assistant. By selecting “Alerts” under “Views” in the far left panel of IT Assistant, administrators can bring up the Alerts page in the right-hand panel. Users can also view traps sent from a particular server blade by selecting the server blade from the IT Assistant system tree and clicking the Log tab. Clicking on a record in the list view opens up a dialogue box that contains detailed trap information.

Running diagnostics

To learn more about an error, administrators can use OpenManage Server Administrator installed on the server blade. Clicking the

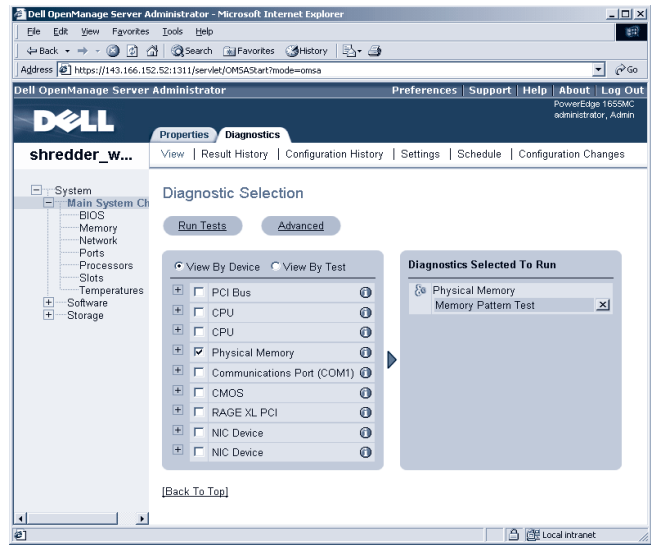


Figure 8. Diagnostic Selection page of OpenManage Server Administrator

The Dell OpenManage software suite provides the tools necessary for administrators to manage the PowerEdge 1655MC server blade system and its modules.

Diagnostics tab from the Server Administrator main window displays the Diagnostic Selection page, where administrators can select “Physical Memory,” for example, and then “Run Tests” (see Figure 8). As the test runs, a progress bar shows the progress of the test.

After the diagnostic test finishes, the program displays the test results. A red “X,” for example, indicates a failed test. Clicking “Details” opens up a message box containing detailed diagnostic results. Users can follow a similar process to investigate errors and run diagnostics on other components of a server blade.

Physically locating a failed server blade

If a server blade fails, administrators might need to locate it physically in a computer room filled with hundreds of racks and thousands of servers. To find the failed blade easily, administrators can run a command using the ERA/MC remote Web console that causes the Status LED on the front cover of the failed server blade to blink. Administrators can start the blinking by following these steps:

1. Identify the service tag of the failed server blade, available from the blade’s Summary and Status pages in IT Assistant (for example, assume the server blade with host name itabld21 has a service tag of BLDIT21)

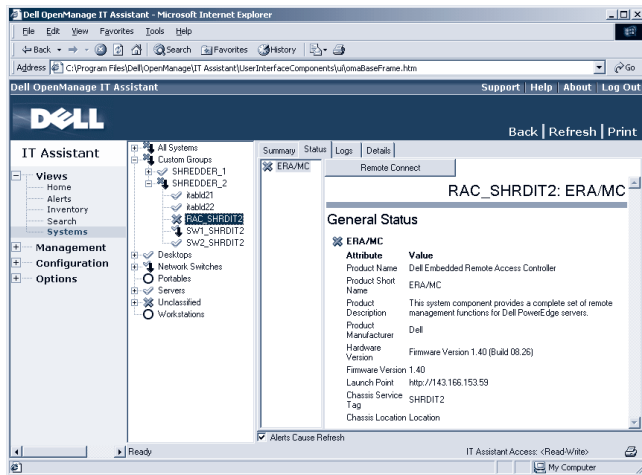


Figure 9. IT Assistant Status page showing ERA/MC error due to integrated Ethernet switch failure

2. From the system tree in IT Assistant, select the ERA/MC (for example, RAC_SHRDIT2) of the enclosure containing the failed server blade
3. Select the Status tab in the right panel
4. Click “Remote Connect” to launch the ERA/MC Web console
5. From the Web console, select the Remote Access tab
6. Find the number assigned to the failed server blade (in this example, BLDIT21 is server-1)
7. Select the Debug tab and select “Command Debug” from the submenu
8. In the Command text box, enter `racadm setled -m server-n ON`, where *n* is the server number found in Step 6

The Status LED on the failed server blade will start blinking, making the blade easy to spot and saving administrators time that may have been wasted by pulling the wrong blade.

Troubleshooting integrated Ethernet switches

Initialization of the integrated Ethernet switch will hang if the POST fails (see the “Errors from integrated Ethernet switches” section earlier in this article). A failed POST makes the integrated switch inaccessible, but the ERA/MC detects the problem and updates its own health status accordingly. In Figure 9, the IT Assistant system tree shows that the integrated switch SW1_SHRDIT2 is powered down and the ERA/MC (RAC_SHRDIT2) is in a critical state, as indicated by the red “X” next to the object.

The IT Assistant Status page provides no further information about what caused the status change. To troubleshoot, administrators can click “Remote Connect” on the ERA/MC Status page from IT Assistant to launch the ERA/MC Web console in a separate browser window (see Figure 10). The Properties tab of the Web console lists the enclosure modules and indicates a critical state by displaying a red “X” next to the module name.

Physically locating a failed integrated Ethernet switch

To find the failed integrated Ethernet switch, administrators must first physically locate the enclosure to which it belongs. The steps are similar to those for locating a failed server blade:

1. From the IT Assistant system tree, select the ERA/MC (for example, RAC_SHRDIT2) of the enclosure to which the failed switch belongs
2. Select the Status tab in the right panel
3. Click “Remote Connect” to launch the ERA/MC Web console
4. Select the Debug tab and select “Command Debug” from the submenu
5. In the Command text box, enter `racadm setled -m chassis ON`

The Status LED on the affected enclosure will start blinking. The integrated switch that has a default host name starting with SW1_ is located on the top of the enclosure, and the integrated switch that has a default host name starting with SW2_ is located on the bottom of the enclosure. After locating the failed switch, administrators can check the Fault Indicator LED on that switch to verify its status. For further troubleshooting, administrators can access the switch through the switch serial console.

Troubleshooting shared enclosure modules

Because the ERA/MC monitors and controls the shared enclosure infrastructure and components, errors on these devices are displayed as a health status change on the ERA/MC. The ERA/MC also sends SNMP traps for these errors. For instance, in Figure 11 the IT Assistant Status page for the ERA/MC shows a warning state for the systems management module, but no further information. To learn more, administrators can check the SNMP traps sent by

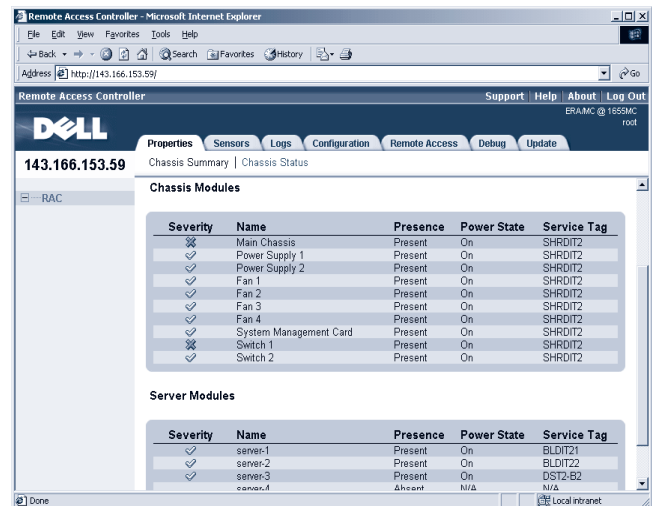


Figure 10. ERA/MC Web console

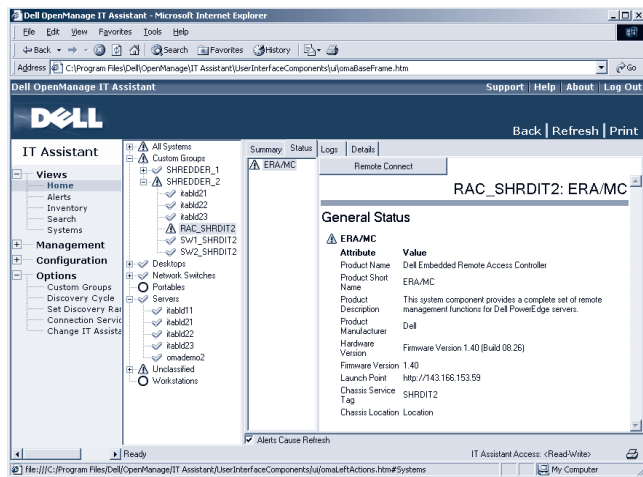


Figure 11. IT Assistant Status page showing warning state of ERA/MC

the ERA/MC; these traps appear in the Alerts page of IT Assistant or the ERA/MC Logs page.

To access the Logs page, administrators can launch the ERA/MC Web console by clicking the Remote Connect button on the ERA/MC Status page of IT Assistant. From the Web console, administrators can select the Logs tab and examine the Embedded System Management (ESM) log for errors. Figure 12 shows an example in which the log suggests that the AC power to Power Supply 2 is lost. By displaying a yellow warning sign in the power supply’s Severity column, the Sensors tab also will show that the AC power for Power Supply 2 is absent.

Physically locating a failed shared enclosure module

Finding a failed shared enclosure module involves the same procedure as that for finding a failed integrated switch:

1. From the IT Assistant system tree, select the ERA/MC (for example, RAC_SHRDIT2) that corresponds to the enclosure containing the failed module
2. Select the Status tab in the right panel
3. Click “Remote Connect” to launch the ERA/MC Web console
4. Select the Debug tab and select “Command Debug” from the submenu
5. In the Command text box, enter `racadm settled -m chassis ON`

The Status LED on the enclosure will start blinking. After locating the enclosure, administrators can check the Fault Indicator LED on each enclosure module for status. They can use the ERA/MC serial command-line interface (CLI) for further troubleshooting.

Modular environments: Lowering costs, boosting performance

Modular computing brings numerous benefits to IT departments, including reduced heat production, reduced power and space

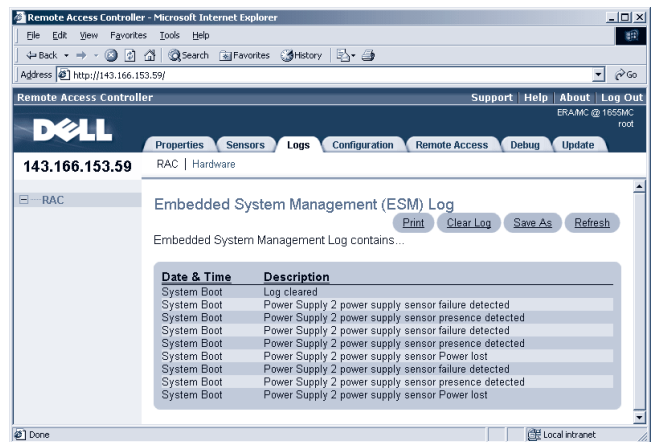


Figure 12. ESM log of the ERA/MC

consumption, and easier deployment. The Dell PowerEdge 1655MC server blade system—the first modular server system released by Dell—can help IT managers to lower data center costs and boost data center performance.

Dell has also enhanced its OpenManage systems management software to support a modular computing environment comprising Dell PowerEdge 1655MC server blade systems. Dell OpenManage IT Assistant, as a centralized management console, provides a single place for IT administrators to check the status of the modular server systems and to monitor error notification from these systems. Additional programs that help manage modular computing environments include OpenManage Server Administrator for server blades, and the Web and serial interfaces of the ERA/MC for integrated Ethernet switches and shared enclosure infrastructures.

All of these management options can help make the PowerEdge 1655MC server blade system easier to manage and troubleshoot, leading to more efficient server administration.

Haihong Zhuo (us_power_solutions@dell.com) is a software engineer in the Dell Product Group—Software Development. She participated in the development of IT Assistant 6.x and is currently on the Systems Management Instrumentation team. Haihong has an M.S. in Computer Engineering from The University of Texas at Austin.

Weimin Pan (us_power_solutions@dell.com) is a senior development engineer in the Dell Remote Management Group. He has worked as a senior system engineer in the Dell Storage Enclosure Subsystem Group. Weimin has an M.S. in Electrical Engineering from the University of Utah and an M.S. in Computer Engineering from Shanghai Jiao Tong University, China.

Dell OpenManage IT Assistant provides a single place for IT administrators to check the status of the modular server systems.