

intel® Technical Advisory

TA-0685-1

5200 NE Elam Young Parkway
Hillsboro, OR 97124

December 10, 2003

Intel® Server Platform SR870BH2 may become unresponsive to IPMI Commands

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice. The Intel® Server Platform SR870BH2 may contain design defects or errors known as errata that may cause the product to deviate from published specifications. Current characterized errata are available on request.

Products Affected

Platforms with the following SKU may be affected: SBHG2

Description

The Server Platform SR870BH2 nodes in a cluster environment may become unresponsive or fail to respond to IPMI commands sent via LAN when the platform is powered down.

Root Cause

The Intel® 82546EB Network Interface Controller EEPROM default configuration allows Address Resolution Protocol responses (ARPs) to route directly to the BMC when the system is in a powered down state (Powered down state = AC source connected and standby current is active). As the volume of ARP traffic increases, the BMC becomes less responsive to IPMI commands sent via the LAN, ultimately ignoring all commands and failing to respond.

The failure thresholds for this issue have been characterized. The thresholds are related to cluster or network scaling and the proportional increase in ARP traffic. The 82546EB Network Interface Controller EEPROM configuration bits can be modified to enable a filter that blocks ARPs that would otherwise be routed to the BMC.

- Disabling the filter allows ARP traffic to route directly to the BMC when the system is in a powered down state. The failure symptoms can be reproduced at 20-50 ARPs / second.
- Enabling the filter routes only the LAN-based IPMI command packets to the BMC and removes ARP traffic to the BMC. The failure symptoms cannot be reproduced at 20-1000 ARPs / second.
- The BMC implementation on the Server Platform SR870BH2 does not utilize ARPs. Enabling the filter does not negatively impact system functionality.

Corrective Action / Resolution

The EEPROM configuration will be modified to enable the ARP filter as the default on all subsequent builds. If the failure symptom is observed, the 82546EB Network Interface Controller EEPROM configuration bits can be modified to enable or disable the filter by running ARPTool64.EFI. This utility is available on IBL.

NOTE: This utility will only modify the EEPROM contents on the Server Platform SR870BH2 with an 82546EB Network Interface Controller. Refer to the "Products Affected" section above to determine if the utility is applicable.

Utility usage and options:

- Copy ARPTool64.EFI to a USB memory device, hard drive, or CD-ROM disk.
- Boot the system to the EFI shell and change to the device containing the utility (i.e. FS0:)
- To **Enable** the Filter: type ARPTool64.EFI **-clear** <Enter> at the EFI prompt
- For **Help** type ARPTool64.EFI **-help** <Enter> at the EFI prompt

Please contact your Intel Support Representative if you require more specific information about this issue.
Enterprise Platforms & Services Division