



US 20030033393A1

(19) **United States**

(12) **Patent Application Publication**

Larson et al.

(10) **Pub. No.: US 2003/0033393 A1**

(43) **Pub. Date: Feb. 13, 2003**

(54) **SYSTEM AND METHOD FOR PROVIDING NETWORK ADDRESS INFORMATION IN A SERVER SYSTEM**

Publication Classification

(51) **Int. Cl.⁷** G06F 15/177; H05K 7/14
(52) **U.S. Cl.** 709/222; 361/797

(76) Inventors: **Thane M. Larson**, Roseville, CA (US);
Kirk Bresniker, Roseville, CA (US)

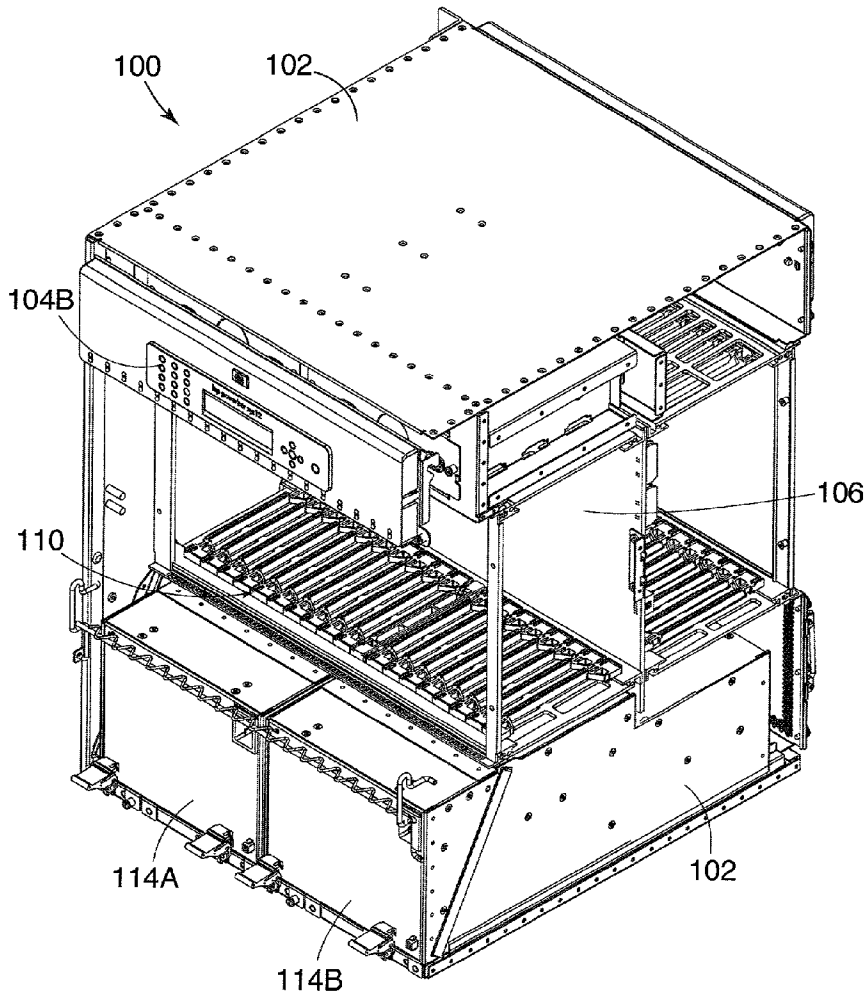
(57) **ABSTRACT**

Correspondence Address:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, CO 80527-2400 (US)

A cPCI server system includes a plurality of host processor cards. A management card is coupled to the plurality of host processor cards via at least one bus. The management card includes at least one user interface for receiving network address information from a user. The management card is configured to send received network address information to the plurality of host processor cards via the at least one bus, thereby configuring the host processor cards for management LAN communications.

(21) Appl. No.: **09/924,029**

(22) Filed: **Aug. 7, 2001**



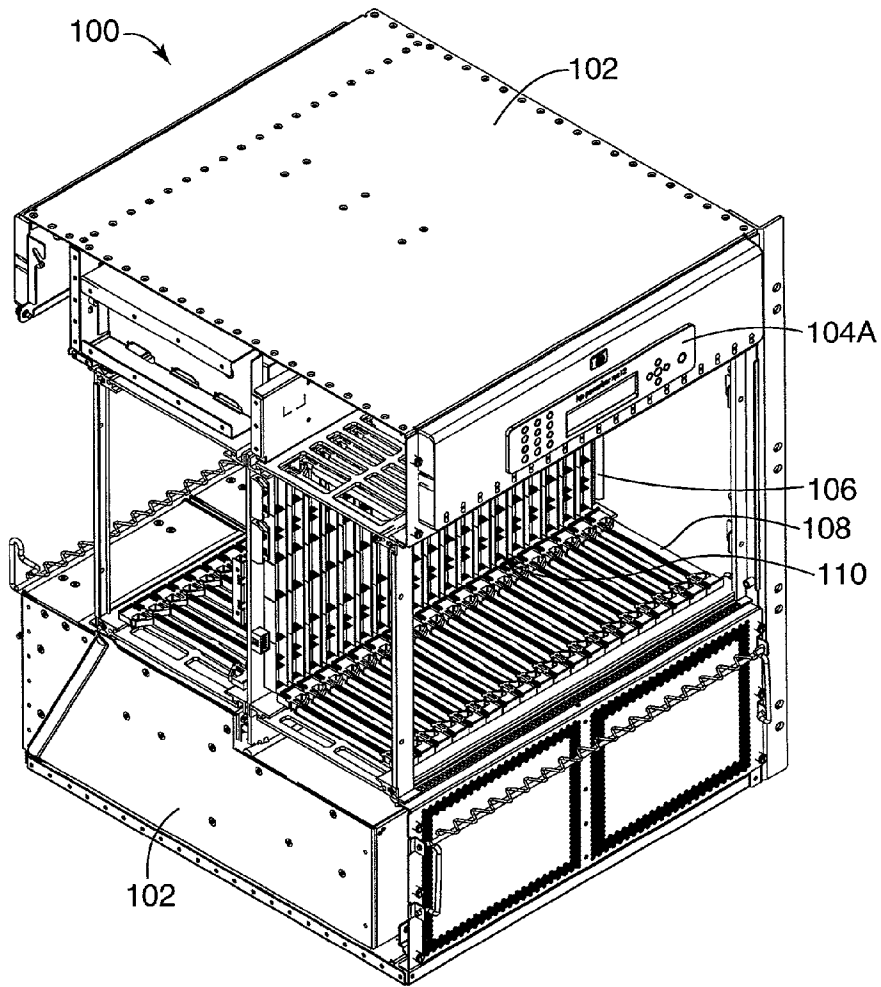


Fig. 1

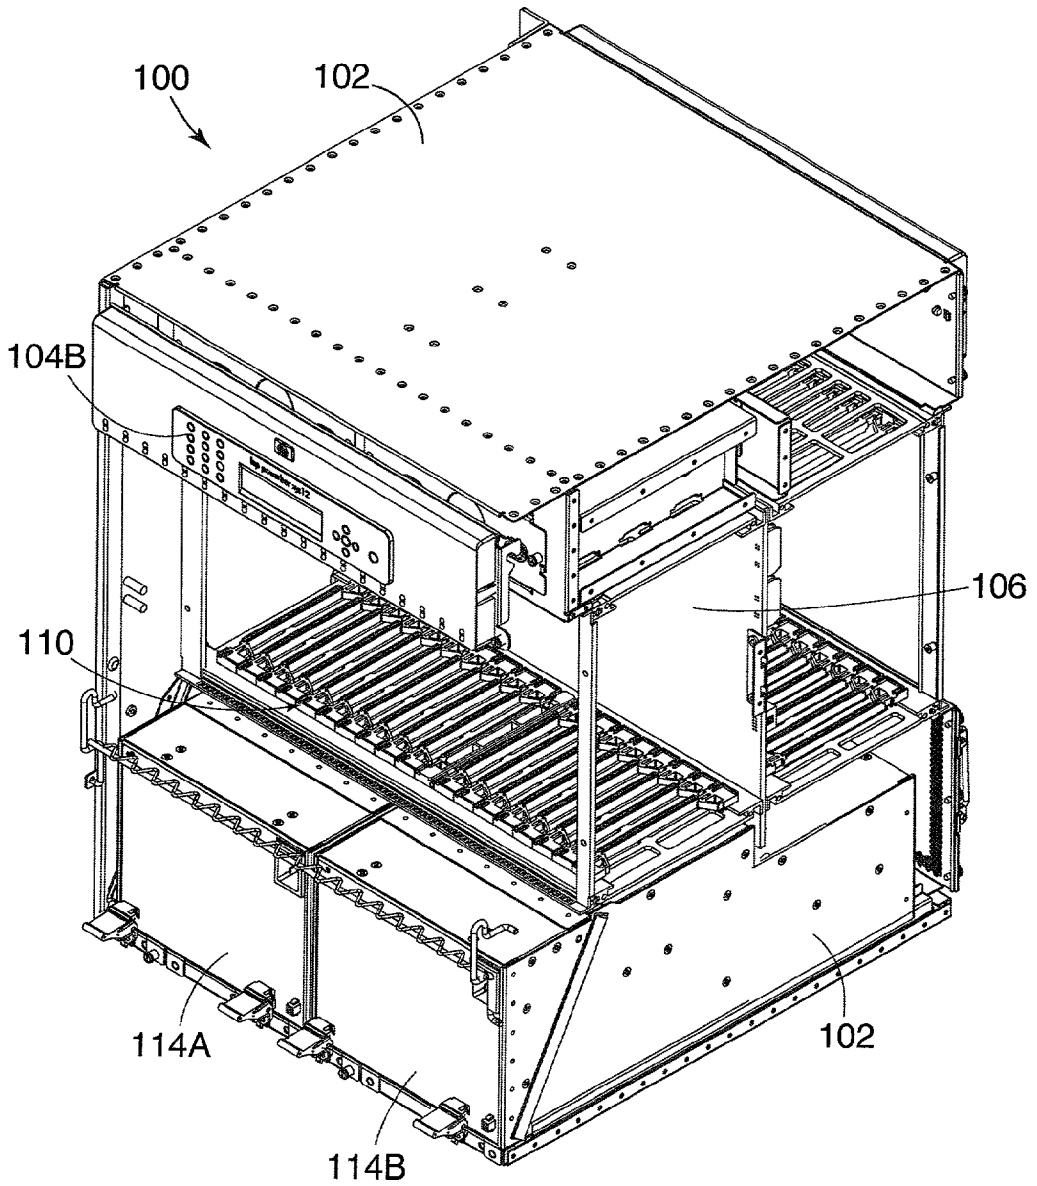


Fig. 2

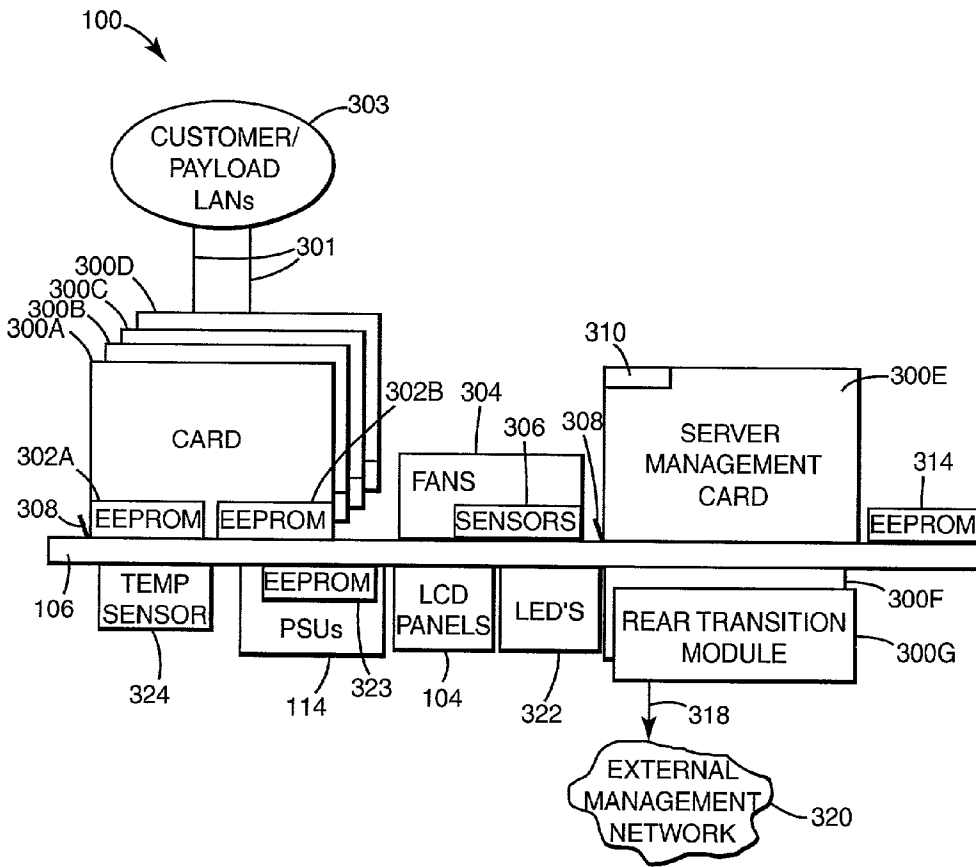


Fig. 3

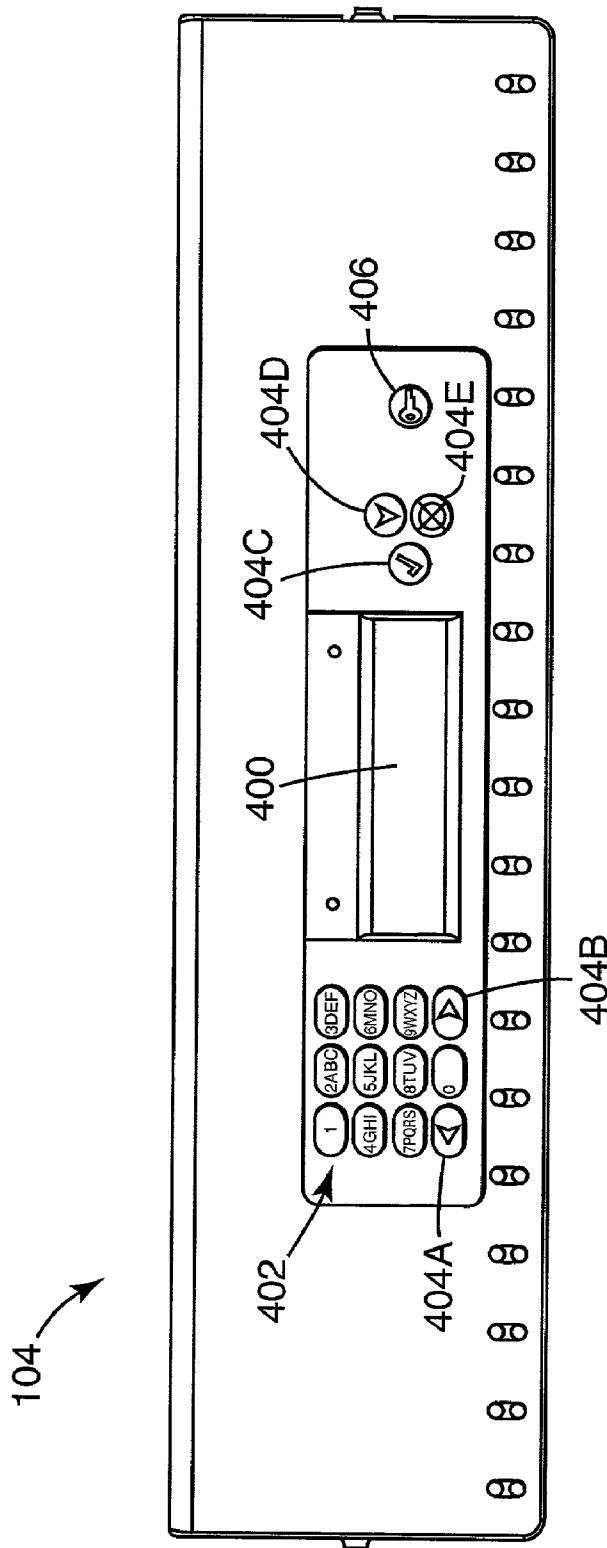


Fig. 4

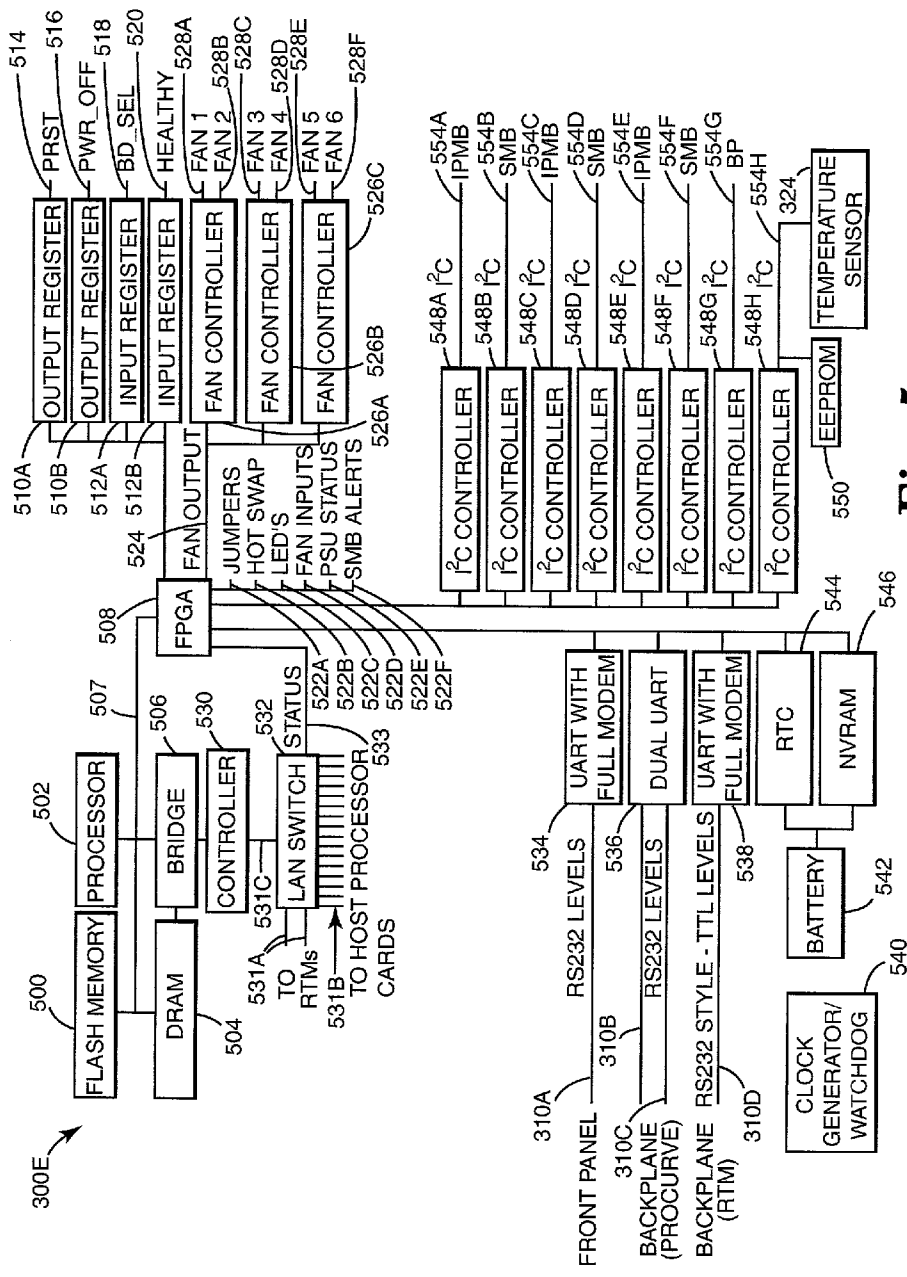


Fig. 5

SYSTEM AND METHOD FOR PROVIDING NETWORK ADDRESS INFORMATION IN A SERVER SYSTEM

THE FIELD OF THE INVENTION

[0001] The present invention relates to server systems. More particularly, the invention relates to a system and method for providing network address information in a cPCI server system for configuration of host processor cards.

BACKGROUND OF THE INVENTION

[0002] Console communications are used in server systems to send and receive status, control and configuration information. If console communications are sent over a LAN using the TCP/IP protocol, the host processor cards on the LAN must be configured and provided Internet protocol (IP) address information. Console LAN communications cannot be established without having the correct IP settings. Most solutions in the marketplace use a serial, RS232 interface as the path to set up the IP address information. This requires a separate port connection to every host processor card. Using an RS-232 port works okay for a standalone server, but when the server is integrated with many others in a common chassis, a separate cable for each server is difficult to maintain and configure. Alternatively, dynamic host configuration protocol (DHCP) could be used to get a random address, but this requires that the proper subnet mask be preset. Furthermore, random IP addresses are not appropriate for servers because clients need to know their exact address.

[0003] It would be desirable to use a single interface in a server system to provide IP address information to all host processor cards in the server system.

SUMMARY OF THE INVENTION

[0004] One form of the present invention provides a server system including a plurality of host processor cards. A management card is coupled to the plurality of host processor cards via at least one bus. The management card includes at least one user interface for receiving network address information from a user. The management card is configured to send received network address information to the plurality of host processor cards via the at least one bus, thereby configuring the host processor cards for management LAN communications.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 is a front perspective view illustrating a server system according to one embodiment of the present invention.

[0006] FIG. 2 is a rear perspective view illustrating the server system shown in FIG. 1.

[0007] FIG. 3 is a block diagram illustrating major components of a server system according to one embodiment of the present invention.

[0008] FIG. 4 is a front view of one of LCD panels used by a server system according to one embodiment of the present invention.

[0009] FIG. 5 is an electrical block diagram illustrating major components of a server management card (SMC) according to one embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0010] In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings that form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural or logical changes may be made without departing from the scope of the present invention. The following detailed description, therefore, is not to be taken in a limiting sense, and the scope of the present invention is defined by the appended claims.

[0011] I. Server System

[0012] FIG. 1 is a front perspective view illustrating a server system 100 according to one embodiment of the present invention. FIG. 2 is a rear perspective view illustrating server system 100. Server system 100 includes panels 102, liquid crystal display (LCD) panels 104A and 104B (collectively referred to as LCD panels 104), backplane 106, chassis 108, and dual redundant power supply units 114A and 114B (collectively referred to as power supply units 114). Panels 102 are attached to chassis 108, and provide protection for the internal components of server system 100. Backplane 106 is positioned near the center of server system 100. Backplane 106 is also referred to as midplane 106. LCD panels 104A and 104B are substantially identical, except for their placement on server system 100. LCD panel 104A is positioned on a front side of server system 100, and LCD panel 104B is positioned on a back side of server system 100.

[0013] Power supply units 114 are positioned at the bottom of server system 100 and extend from a back side of server system 100 to a front side of server system 100. Power supply units 114 each include an associated cooling fan 304 (shown in block form in FIG. 3). In one form of the invention, additional cooling fans 304 are positioned behind LCD panel 104B. In one embodiment, 4 chassis cooling fans 304 are used in server system 100. In an alternative embodiment, 6 chassis cooling fans 304 are used. Other numbers and placement of cooling fans 304 may be used. In one form of the invention, cooling fans 304 form an N+1 redundant cooling system, where "N" represents the total number of necessary fans 304, and "1" represents the number of redundant fans 304.

[0014] In one embodiment, server system 100 supports the Compact Peripheral Component Interconnect (cPCI) form factor of printed circuit assemblies (PCAs). Server system 100 includes a plurality of cPCI slots 110 for receiving cards/modules 300 (shown in block form in FIG. 3). In one embodiment, system 100 includes ten slots 110 on each side of backplane 106 (referred to as the 10 slot configuration). In an alternative embodiment, system 100 includes nineteen slots 110 on each side of backplane 106 (referred to as the 19 slot configuration). Additional alternative embodiments use other slot configurations.

[0015] FIG. 3 is a block diagram illustrating major components of server system 100. Server system 100 includes backplane 106, a plurality of cards/modules 300A-300G (collectively referred to as cards 300), fans 304, electrically erasable programmable read only memory (EEPROM) 314, LEDs 322, LCD panels 104, power supply units (PSUs) 114,

and temperature sensor 324. Cards 300 are inserted in slots 110 (shown in FIGS. 1 and 2) in system 100. In one form of the invention, cards 300 may occupy more than one slot 110. In one embodiment, cards 300 include host processor cards 300A, hard disk cards 300B, managed Ethernet switch cards 300C and 300D, a server management card (SMC) 300E, and 2 redundant SMC local area network (LAN) rear transition modules (RTMs) 300F and 300G. In one embodiment, there is one managed Ethernet switch card 300C fitted in the 10 slot chassis embodiment, and up to two managed Ethernet switch cards 300C and 300D fitted in the 19 slot chassis embodiment. In one form of the invention, managed Ethernet switch cards 300C and 300D are "Procurve" managed Ethernet switch cards.

[0016] In one embodiment, two types of host processor cards 300A may be used in server system 100—PA-RISC host processor cards and IA32 host processor cards. Multiple host processor cards 300A and hard disk cards 300B are used in embodiments of server system 100, but are each represented by a single card in FIG. 3 to simply the figure. In one form of the invention, up to 8 host processor cards 300A are used in the 10 slot configuration, and up to 16 host processor cards 300A are used in the 19 slot configuration. In one embodiment, each of cards 300 can be hot swapped.

[0017] In one embodiment, cards 300 each include a pair of EEPROMs 302A and 302B, which are discussed below. Power supply units 114 each include an EEPROM 323 for storing power supply identification and status information. Fans 304 include associated sensors 306 for monitoring the speed of the fans 304. In one embodiment, LEDs 322 include eight status LEDs, six LAN LEDs to indicate the speed and link status of LAN links 318, a blue hot swap status LED to indicate the ability to hot swap SMC 300E, a power-on indicator LED, and three fan control indicator LEDs.

[0018] The operational health of cards 300 and system 100 are monitored by SMC 300E to ensure the reliable operation of the system 100. SMC 300E includes serial ports 310 (discussed below), and an extraction lever 308 with an associated switch. In one embodiment, all cards 300 include an extraction lever 308 with an associated switch.

[0019] In one embodiment, SMC 300E is the size of a typical compact PCI (cPCI) card, and supports PA-RISC and the IA32 host processor cards 300A. SMC 300E electrically connects to other components in system 100, including cards 300, temperature sensor 324, power supply units 114, fans 304, EEPROM 314, LCD panels 104, LEDs 322, and SMC rear transition modules 300F and 300G via backplane 106. In most cases, the connections are via I²C buses 554 (shown in FIG. 5), as described in further detail below. The I²C buses 554 allow bi-directional communication so that status information can be sent to SMC 300E and configuration information sent from SMC 300E. In one embodiment, SMC 300E uses I²C buses 554 to obtain environmental information from power supply units 114, host processor cards 300A, and other cards 300 fitted into system 100.

[0020] SMC 300E also includes a LAN switch 532 (shown in FIG. 5) to connect console management LAN signals from the host processor cards 300A to an external management network (also referred to as management LAN) 320 via one of the two SMC rear transition modules 300F and 300G. In one embodiment, the two SMC rear transition

modules 300F and 300G each provide external 10/100Base-T LAN links 318 for connectivity to management LAN 320. In one embodiment, SMC rear transition modules 300F and 300G are fibre-channel, port-bypass cards

[0021] Managed Ethernet switch cards 300C and 300D are connected to host processor cards 300A through backplane 106, and include external 10/100/1000Base-T LAN links 301 for connecting host processor cards to external customer or payload LANs 303. managed Ethernet switch cards 300C and 300D are fully managed LAN switches.

[0022] II. LCD Panel

[0023] FIG. 4 is a front view of one of LCD panels 104. In one form of the invention, each LCD panel 104 includes a 2x20 LCD display 400, 10 alphanumeric keys 402, 5 menu navigation/activation keys 404A-404E (collectively referred to as navigation keys 404), and a lockout key 406 with associated LED (not shown) that lights lockout key 406. If a user presses a key 402, 404, or 406, an alert signal is generated and SMC 300E polls the LCD panels 104A and 104B to determine which LCD panel was used, and the key that was pressed.

[0024] Alphanumeric keys 402 allow a user to enter alphanumeric strings that are sent to SMC 300E. Navigation keys 404 allow a user to navigate through menus displayed on LCD display 400, and select desired menu items. Navigation keys 404A and 404B are used to move left and right, respectively, within the alphanumeric strings. Navigation key 404C is an "OK/Enter" key. Navigation key 404D is used to move down. Navigation key 404E is a "Cancel" key.

[0025] LCD panels 104 provide access to a test shell (discussed below) that provides system information and allows configuration of system 100. As discussed below, other methods of access to the test shell are also provided by system 100. To avoid contention problems between the two LCD panels 104, and the other methods of access to the test shell, a lockout key 406 is provided on LCD panels 104. A user can press lockout key 406 to gain or release control of the test shell. In one embodiment, lockout key 406 includes an associated LED to light lockout key 406 and indicate a current lockout status.

[0026] In one embodiment, LCD panels 104 also provide additional information to that displayed by LEDs 322 during start-up. If errors are encountered during the start-up sequence, LCD panels 104 provide more information about the error without the operator having to attach a terminal to one of the SMC serial ports 310.

[0027] III. Server Management Card (SMC)

[0028] A. SMC Overview

[0029] FIG. 5 is an electrical block diagram illustrating major components of server management card (SMC) 300E. SMC 300E includes flash memory 500, processor 502, dynamic random access memory (DRAM) 504, PCI bridge 506, field programmable gate array (FPGA) 508, output registers 510A and 510B, input registers 512A and 512B, fan controllers 526A-526C (collectively referred to as fan controllers 526), network controller 530, LAN switch 532, universal asynchronous receiver transmitter (UART) with modem 534, dual UART 536, UART with modem 538, clock generator/watchdog 540, battery 542, real time clock (RTC)

544, non-volatile random access memory (NVRAM) **546**, I²C controllers **548A-548H** (collectively referred to as I²C controllers **548**), EEPROM **550**, and temperature sensor **324**. In one embodiment, components of SMC **300E** are connected together via PCI buses **507**. In one form of the invention, PCI buses **507** are not routed between slots **110**. Switched LAN signals through LAN switch **532** are routed between slots **110**.

[**0030**] Functions of SMC **300E** include supervising the operation of other components within system **100** (e.g. fan speed, temperature, card present) and reporting their health to a central location (e.g., external management network **320**), reporting any failures to a central location (e.g., external management network **320**), providing a LAN switch **532** to connect console management LAN signals from the SMC **300E** and host processor cards **300A** to an external management network **320**, and providing an initial boot configuration for the system **100**.

[**0031**] B. SMC Processor And Memory

[**0032**] SMC **300E** includes chassis management processor **502**. In one embodiment, chassis management processor **502**, also referred to as SMC processor **502**, is a StrongARM SA-110 processor with supporting buffer. In one embodiment, SMC **300E** uses a Linux operating system. SMC **300E** also runs server management application (SMA) software/firmware. In one embodiment, the operating system and SMA are stored in flash memory **500**. In one form of the invention, all information needed to power-up SMC **300E**, and for SMC **300E** to become operational, are stored in flash memory **500**. In one embodiment, flash memory **500** includes 4 to 16 Mbytes of storage space to allow SMC **300E** to boot-up as a stand-alone card (i.e., no network connection needed).

[**0033**] SMC **300E** also includes DRAM **504**. In one embodiment, DRAM **504** includes 32, 64 or 128 Mbytes of storage space. In one form of the invention, a hardware fitted table is stored in DRAM **504**. The hardware fitted table includes information representing the physical configuration of system **100**. The hardware fitted table changes if there is a physical change to system **100**, such as by a hardware device being added to or removed from system **100**. The hardware fitted table includes hardware type information (e.g., whether a device is an IA32/PA-RISC/Disk Carrier/RTM (i.e., rear transition module)/PSU/LCD panel/Modem/Unknown device, etc.), hardware revision and serial number, status information, configuration information, and hot-swap status information.

[**0034**] Processor **502** is coupled to FPGA **508**. FPGA **508** includes 6 sets of input/output lines **522A-522F**. Lines **522A** are connected to jumpers for configuring SMC **300E**. Lines **522B** are hot swap lines for monitoring the hot swap status of cards **300**. In one embodiment, hot swap lines **522B** include **18** hot swap status input lines, which allow SMC **300E** to determine the hot swap status of the host processor cards **300A**, hard disk cards **300B**, managed Ethernet switch cards **300C** and **300D**, SMC rear transition modules **300F** and **300G**, and power supply units **114**. Lines **522C** are LED lines that are coupled to LEDs **322**. Lines **522D** are fan input lines that are coupled to fan sensors **306** for monitoring the speed of fans **304**. Lines **522E** are power supply status lines that are coupled to power supply units **114** for determining whether both, or only one power supply unit **114** is present.

Lines **522F** are SMB alert lines for communicating alert signals related to SMB I²C buses **554B**, **554D**, and **554F**.

[**0035**] C. Clock, Battery & NVRAM

[**0036**] SMC **300E** includes a real time clock (RTC) **544** and an associated battery **542** to preserve the clock. Real time clock **544** provides the correct time of day. SMC **300E** also includes NVRAM **546** for storing clock information. In one embodiment, NVRAM **546** uses the same battery as real time clock **544**.

[**0037**] D. LAN switch

[**0038**] SMC **300E** sends and receives management LAN communications through PCI bridge **506** and controller **530** to LAN switch **532**. In one embodiment, LAN switch **532** is an unmanaged LAN switch including **19** ports, with two ports connected to SMC rear transition modules **300F** and **300G** (shown in FIG. 3) via links **531A** for communications with external management network **320** (shown in FIG. 3), **16** ports for connecting to the management LAN connections of up to 16 host processor cards **300A** via links **531B** through backplane **106**, and one port for connecting to the SMC's LAN port (i.e., output of controller **530**) via links **531C**. SMC **300E** provides management support for console LAN management signals sent and received through LAN switch **532**. SMC **300E** provides control of management LAN signals of host processor cards **300A**, managed Ethernet switches **300C** and **300D**, SMC processor **502**, and SMC rear transition modules **300F** and **300G**. SMC **300E** monitors the status of the management LAN connections of up to 16 host processor cards **300A** to LAN switch **532**, and reports an alarm event if any of the connections are lost. FPGA **508** and LAN switch **532** are coupled together via an RS-232 link **533** for the exchange of control and status information.

[**0039**] E. I²C Buses

[**0040**] Server system **100** includes eight I²C buses **554A-554H** (collectively referred to as I²C buses **554**) to allow communication with components within system **100**. I²C buses **554** are coupled to FPGA **508** via I²C controllers **548**. In one embodiment, the I²C buses **554** include 3 intelligent platform management bus (IPMB) buses **554A**, **554C**, and **554E**, 3 system management bus (SMB) buses **554B**, **554D**, and **554F**, a backplane ID bus (BP) **554G**, and an I²C bus **554H** for accessing SMC EEPROM **550** and chassis temperature sensor **324**. A different number and configuration of I²C buses **554** may be used depending upon the desired implementation. SMC **300E** maintains a system event log (SEL) within non-volatile flash memory **500** for storing information gathered over I²C buses **554**.

[**0041**] The IPMB I²C buses **554A**, **554C**, and **554E** implement the intelligent platform management interface (IPMI) specification. The IPMI specification is a standard defining an abstracted interface to platform management hardware. IPMI is layered over the standard I²C protocol. SMC **300E** uses one or more of the IPMB I²C buses **554A**, **554C**, and **554E** to retrieve static data from each of the host processor cards **300A** and hard disk cards **300B**. The static data includes identification information for identifying each of the cards **300A** and **300B**. Each slot **110** in system **100** can be individually addressed to retrieve the static configuration data for the card **300** in that slot **110**. In one embodiment, the host processor cards **300A** and hard disk cards **300B** each

include an EEPROM 302A (shown in FIG. 3) that stores the static identification information retrieved over IPMB I²C buses 554A, 554C, and 554E. In one embodiment, each EEPROM 302A contains the type of card, the name of the card, the hardware revision of the card, the card's serial number and card manufacturing information.

[0042] SMC 300E also uses one or more of the IPMB I²C buses 554A, 554C, and 554E, to retrieve dynamic environmental information from each of the host processor cards 300A and hard disk cards 300B. In one embodiment, this dynamic information is held in a second EEPROM 302B (shown in FIG. 3) on each of the cards 300A and 300B. In one form of the invention, the dynamic board data includes card temperature and voltage measurements. In one embodiment, SMC 300E can write information to the EEPROMs 302A and 302B on cards 300.

[0043] The three SMB I²C buses 554B, 554D, and 554F also implement the IPMI specification. The three SMB I²C buses 554B, 554D, and 554F, are coupled to LEDs 322, the two LCD panels 104, the dual redundant power supply units 114, and some of the host processor cards 300A. SMC 300E uses one or more of the SMB I²C buses 554B, 554D, and 554F, to provide console communications via the LCD panels 104. In order for the keypad key-presses on the LCD panels 104 to be communicated back to SMC 300E, an alert signal is provided when keys are pressed that causes SMC 300E to query LCD panels 104 for the keys that were pressed.

[0044] SMC 300E communicates with power supply units 114 via one or more of the SMB I²C buses 554B, 554D, and 554F to obtain configuration and status information including the operational state of the power supply units 114. In one embodiment, the dual redundant power supply units 114 provide voltage rail measurements to SMC 300E. A minimum and maximum voltage value is stored by the power supply units 114 for each measured rail. The voltage values are polled by SMC 300E at a time interval defined by the current configuration information for SMC 300E. If a voltage measurement goes out of specification, defined by maximum and minimum voltage configuration parameters, SMC 300E generates an alarm event. In one embodiment, power supply units 114 store configuration and status information in their associated EEPROMs 323 (shown in FIG. 3).

[0045] Backplane ID Bus (BP) 554G is coupled to backplane EEPROM 314 (shown in FIG. 3) on backplane 106. SMC 300E communicates with the backplane EEPROM 314 over the BP bus 554G to obtain backplane manufacturing data, including hardware identification and revision number. On start-up, SMC 300E communicates with EEPROM 314 to obtain the manufacturing data, which is then added to the hardware fitted table. The manufacturing data allows SMC 300E to determine if it is in the correct chassis for the configuration it has on board, since it is possible that the SMC 300E has been taken from a different chassis and either hot-swapped into a new chassis, or added to a new chassis and the chassis is then powered up. If there is no valid configuration on board, or SMC 300E cannot determine which chassis it is in, then SMC 300E waits for a pushed configuration from external management network 320, or for a manual user configuration via one of the connection methods discussed below.

[0046] In one embodiment, there is a single temperature sensor 324 within system 100. SMC 300E receives temperature information from temperature sensor 324 over I²C bus 554H. SMC 300E monitors and records this temperature and adjusts the speed of the cooling fans 304 accordingly, as described below. SMC also uses I²C bus 554H to access EEPROM 550, which stores board revision and manufacture data for SMC 300E.

[0047] F. Serial Ports

[0048] SMC 300E includes 4 RS-232 interfaces 310A-310D (collectively referred to as serial ports 310). RS-232 serial interface 310A is via a 9-pin Male D-type connector on the front panel of SMC 300E. The other three serial ports 310B-310D are routed through backplane 106. The front panel RS-232 serial interface 310A is connected via a UART with a full modem 534 to FPGA 508, to allow monitor and debug information to be made available via the front panel of SMC 300E. Backplane serial port 310D is also connected via a UART with a full modem 538 to FPGA 508. In one embodiment, backplane serial port 310D is intended as a debug or console port. The other two backplane serial interfaces 310B and 310C are connected via a dual UART 536 to FPGA 508, and are routed to managed Ethernet switches 300C and 300D through backplane 106. These two backplane serial interfaces 310B and 310C are used to connect to and configure the managed Ethernet switch cards 300C and 300D, and to obtain status information from the managed Ethernet switch cards 300C and 300D.

[0049] G. Fans And Temperature Control

[0050] In one embodiment, server system 100 includes six chassis fans 304. Server system 100 includes temperature sensor 324 to monitor the chassis temperature, and fan sensors 306 to monitor the six fans 304. In one embodiment, fan sensors 306 indicate whether a fan 304 is rotating and the fan's speed setting. In one form of the invention, FPGA 508 includes 6 fan input lines 522D (i.e., one fan input line 522D from each fan sensor 306) to monitor the rotation of the six fans 304, and a single fan output line 524 coupled to fan controllers 526A-526C. Fan controllers 526A-526C control the speed of fans 304 by a PWM (pulse width modulation) signal via output lines 528A-528F. If a fan 304 stalls, the monitor line 522D of that fan 304 indicates this condition to FPGA 508, and an alarm event is generated. The speed of fans 304 is varied to maintain an optimum operating temperature versus fan noise within system 100. If the chassis temperature sensed by temperature sensor 324 reaches or exceeds a temperature alarm threshold, an alarm event is generated. When the temperature reduces below the alarm threshold, the alarm event is cleared. If the temperature reaches or exceeds a temperature critical threshold, the physical integrity of the components within system 100 are considered to be at risk, and SMC 300E performs a system shut-down, and all cards 300 are powered down except SMC 300E. When the chassis temperature falls below the critical threshold and has reached the alarm threshold, SMC 300E restores the power to all of the cards 300 that were powered down when the critical threshold was reached.

[0051] In one embodiment, SMC 300E controls the power state of cards 300 using power reset (PRST) lines 514 and power off (PWR_OFF) lines 516. FPGA 508 is coupled to power reset lines 514 and power off lines 516 via output registers 510A and 510B, respectively. In one embodiment,

power reset lines **514** and power off lines **516** each include 19 output lines that are coupled to cards **300**. SMC **300E** uses power off lines **516** to turn off the power to selected cards **300**, and uses power reset lines **514** to reset selected cards **300**. In one embodiment, a lesser number of power reset and power off lines are used for the 10 slot chassis configuration.

[0052] H. Clock Generator/Watchdog

[0053] SMC **300E** is protected by both software and hardware watchdog timers. The watchdog timers are part of clock generator/watchdog block **540**, which also provides a clock signal for SMC **300E**. The hardware watchdog timer is started before software loading commences to protect against failure. In one embodiment, the time interval is set long enough to allow a worst-case load to complete. If the hardware watchdog timer expires, SMC processor **502** is reset.

[0054] I. Modes Of Operation

[0055] In one embodiment, SMC **300E** has three phases or modes of operation—Start-up, normal operation, and hot swap. The start-up mode is entered on power-up or reset, and controls the sequence needed to make SMC **300E** operational. SMC **300E** also provides minimal configuration information to allow chassis components to communicate on the management LAN. The progress of the start-up procedure can be followed on LEDs **322**, which also indicate any errors during start-up.

[0056] The normal operation mode is entered after the start-up mode has completed. In the normal operation mode, SMC **300E** monitors the health of system **100** and its components, and reports alarm events. SMC **300E** monitors the chassis environment, including temperature, fans, input signals, and the operational state of the host processor cards **300A**.

[0057] SMC **300E** reports alarm events to a central point, namely an alarm event manager, via the management LAN (i.e., through LAN switch **532** and one of the two SMC rear transition modules **300F** or **300G** to external management network **320**). The alarm event manager is an external module that is part of external management network **320**, and that handles the alarm events generated by server system **100**. The alarm event manager decides what to do with received alarms and events, and initiates any recovery or reconfiguration that may be needed. In addition to sending the alarm events across the management network, a system event log (SEL) is maintained in SMC **300E** to keep a record of the alarms and events. The SEL is held in non-volatile flash memory **500** in SMC **300E** and is maintained over power cycles, and resets of SMC **300E**.

[0058] In the normal operation mode, SMC **300E** may receive and initiate configuration commands and take action on received commands. The configuration commands allow the firmware of SMC processor **502** and the hardware controlled by processor **502** to be configured. This allows the operation of SMC **300E** to be customized to the current environment. Configuration commands may originate from the management network **320**, one of the local serial ports **310** via a test shell (discussed below), or one of the LCD panels **104**.

[0059] The hot swap mode is entered when there is an attempt to remove a card **300** from system **100**. In one

embodiment, all of the chassis cards **300** can be hot swapped, including SMC **300E**, and the two power supply units **114**. An application shutdown sequence is initiated if a card **300** is to be removed. The shutdown sequence performs all of the steps needed to ready the card **300** for removal.

[0060] In one embodiment, FPGA **508** includes **18** hot swap status inputs **522B**. These inputs **522B** allow SMC **300E** to determine the hot swap status of host processor cards **300A**, hard disk cards **300B**, managed Ethernet switch cards **300C** and **300D**, SMC rear transition module cards **300F** and **300G**, and power supply units **114**. The hot-swap status of the SMC card **300E** itself is also determined through this interface **522B**.

[0061] An interrupt is generated and passed to SMC processor **502** if any of the cards **300** in system **100** are being removed or installed. SMC **300E** monitors board select (BD_SEL) lines **518** and board healthy (HEALTHY) lines **520** of cards **300** in system **100**. In one embodiment, board select lines **518** and healthy lines **520** each include 19 input lines, which are connected to FPGA **508** via input registers **512A** and **512B**, respectively. SMC **300E** monitors the board select lines **518** to sense when a card **300** is installed. SMC **300E** monitors the healthy lines **520** to determine whether cards **300** are healthy and capable of being brought out of a reset state.

[0062] When SMC **300E** detects that a card has been inserted or removed, an alarm event is generated. When a new card **300** is inserted in system **100**, SMC **300E** determines the type of card **300** that was inserted by polling the identification EEPROM **302A** of the card **300**. Information is retrieved from the EEPROM **302A** and added to the hardware fitted table. SMC **300E** also configures the new card **300** if it has not been configured, or if its configuration differs from the expected configuration. When a card **300**, other than the SMC **300E**, is hot-swapped out of system **100**, SMC **300E** updates the hardware fitted table accordingly.

[0063] In one embodiment, SMC **300E** is extracted in three stages: (1) an interrupt is generated and passed to the SMC processor **502** when the extraction lever **308** on the SMC front panel is set to the “extraction” position in accordance with the Compact PCI specification, indicating that SMC **300E** is about to be removed; (2) SMC processor **502** warns the external management network **320** of the SMC **300E** removal and makes the extraction safe; and (3) SMC processor **502** indicates that SMC may be removed via the blue hot swap LED **322**. SMC **300E** ensures that any application download and flashing operations are complete before the hot swap LED **322** indicates that the card **300E** may be removed.

[0064] J. User Connectivity

[0065] In one embodiment, there are two test shells implemented within SMC **300E**. There is an application level test shell that is a normal, run-time, test shell accessed and used by users and applications. There is also a stand-alone test shell that is a manufacturer test shell residing in flash memory **500** that provides manufacturing level diagnostics and functions. The stand-alone test shell is activated when SMC **300E** boots and an appropriate jumper is in place on SMC **300E**. The stand-alone test shell allows access to commands that the user would not, or should not have access to.

[0066] The test shells provide an operator interface to SMC 300E. This allows an operator to query the status of system 100 and (with the required authority level) to change the configuration of system 100.

[0067] A user can interact with the test shells by a number of different methods, including locally via a terminal directly attached to one of the serial ports 310, locally via a terminal attached by a modem to one of the serial ports 310, locally via one of the two LCD panels 104, and remotely via a telnet session established through the management LAN 320. A user may connect to the test shells by connecting a terminal to either the front panel serial port 31 OA or rear panel serial ports 310B-310D of SMC 300E, depending on the console/modem serial port configuration. The RS-232 and LAN connections provide a telnet console interface. LCD panels 104 provide the same command features as the telnet console interface. SMC 300E can function as either a dial-in facility, where a user may establish a link by calling to the modem, or as a dial-out facility, where SMC 300E can dial out to a configured number.

[0068] The test shells provide direct access to alarm and event status information. In addition, the test shells provides the user with access to other information, including temperature logs, voltage logs, chassis card fitted table, and the current setting of all the configuration parameters. The configuration of SMC 300E may be changed via the test shells. Any change in configuration is communicated to the relevant cards 300 in system 100. In one embodiment, configuration information downloaded via a test shell includes a list of the cards 300 expected to be present in system 100, and configuration data for these cards 300. The configuration information is stored in flash memory 500, and is used every time SMC 300E is powered up.

[0069] K. Console LAN Configuration

[0070] In one embodiment, SMC 300E uses the I²C buses 554 coupled to host processor cards 300A to set the internet protocol (IP) address information. As mentioned above, the I²C buses 554 coupled to host processor cards 300A use the intelligent platform management interface (IPMI). IPMI provides the definition to the structure of basic I²C master read and write commands. In one form of the invention, SMC 300E augments the IPMI protocol with additional host processor card configuration commands that send over the IPMI I²C buses 554 the IP information required to get a console or management LAN up and running. In one embodiment, the configuration parameters downloaded from SMC 300E to host processor cards 300A over IPMI I²C buses 554 include: IP address, gateway address, subnet address, and host name. A host processor card 300A that needs the use of a console LAN, will always have its IPMI interface active, despite the boot state of the host processor card 300A, including powered off. With the possibility that a customer could move a working host processor card 300A to a different chassis and possibly a different LAN, what was previously a unique LAN address could later conflict with an existing LAN address. So the ability to set LAN address information in an inoperable state of a host processor card 300A is important.

[0071] By using the already existing IPMI I²C buses 554 to set the IP address information, there is virtually no added cost to configure system 100, and system 100 is robust. In contrast to the prior art technique of providing an RS-232

connection to each individual server to set IP address information, the use of an IPMI I²C bus 554 allows for multiple devices on one bus, thus greatly reducing the complexity of the system. This also allows for better remote management, since it is undesirable to have multiple connections into a server, or having to physically be present at the server to change configurations. A user may log in remotely to management card 300E through management LAN 320, provide the IP address information to management card 300E, which then sends the IP address information to host processor cards 300A using the IPMI I²C buses 554.

[0072] Although specific embodiments have been illustrated and described herein for purposes of description of the preferred embodiment, it will be appreciated by those of ordinary skill in the art that a wide variety of alternate and/or equivalent implementations may be substituted for the specific embodiments shown and described without departing from the scope of the present invention. Those with skill in the chemical, mechanical, electro-mechanical, electrical, and computer arts will readily appreciate that the present invention may be implemented in a very wide variety of embodiments. This application is intended to cover any adaptations or variations of the preferred embodiments discussed herein. Therefore, it is manifestly intended that this invention be limited only by the claims and the equivalents thereof.

What is claimed is:

1. A server system comprising:

a plurality of host processor cards;

a management card coupled to the plurality of host processor cards via at least one bus, the management card including at least one user interface for receiving network address information from a user, the management card configured to send received network address information to the plurality of host processor cards via the at least one bus, thereby configuring the host processor cards for management LAN communications.

2. The server system of claim 1, wherein the at least one bus is an I²C bus.

3. The server system of claim 2, wherein the at least one bus is an intelligent platform management interface (IPMI) I²C bus.

4. The server system of claim 3, wherein the network address information sent from the management card to the plurality of host processor cards is sent using an augmented IPMI protocol that includes additional host processor card configuration commands.

5. The server system of claim 1, wherein the network address information includes internet protocol (IP) address information.

6. The server system of claim 5, wherein the IP address information includes an IP address, gateway address, subnet address, and host name.

7. The server system of claim 1, wherein the at least one user interface includes at least one serial port and at least one LAN interface.

8. The server system of claim 7, wherein the at least one user interface further includes at least one LCD panel mounted on the server system.

9. A server management card for a server system having a plurality of host processor cards, the server management card comprising:

at least one user interface for allowing a user to enter network address information;

at least one I²C bus connection for connecting the server management card to the plurality of host processor cards via at least one I²C bus; and

a controller configured to output entered network address information to the plurality of host processor cards via the at least one I²C bus connection, thereby configuring the plurality of host processor cards for network communications.

10. The server management card of claim 9, wherein the at least one I²C bus is an intelligent platform management interface (IPMI) I²C bus.

11. The server management card of claim 10, wherein the network address information output from the server management card to the plurality of host processor cards is sent using an augmented IPMI protocol that includes additional host processor card configuration commands.

12. The server management card of claim 9, wherein the network address information includes internet protocol (IP) address information.

13. The server management card of claim 12, wherein the IP address information includes an IP address, gateway address, subnet address, and host name.

14. The server management card of claim 9, wherein the at least one user interface includes at least one serial port, at least one LAN interface, and at least one LCD panel mounted on the server system.

15. A method of configuring host processor cards in a server system for management network communications, the method comprising:

providing a management card in the server system having at least one user interface;

providing at least one bus connecting the management card and the host processor cards;

entering network address information to the management card through the at least one user interface; and

sending entered network address information from the management card to the host processor cards, thereby configuring the host processor cards for management network communications.

16. The method of claim 15, wherein the at least one bus is an I²C bus.

17. The method of claim 16, wherein the at least one bus is an intelligent platform management interface (IPMI) I²C bus.

18. The method of claim 17, wherein the network address information sent from the management card to the host processor cards is sent using an augmented IPMI protocol that includes additional host processor card configuration commands.

19. The method of claim 15, wherein the network address information includes an internet protocol (IP) address, gateway address, subnet address, and host name.

20. The method of claim 15, wherein the at least one user interface includes at least two of a serial port, a LAN interface, and an LCD panel mounted on the server system.

* * * * *